THE PERFORMANCE OF HIGHLY ACTIVE PROBLEM SOLVING STRATEGIES IN NOVEL TASK ENVIRONMENTS

by

Gary Scott Mahon

A Dissertation Submitted to the Faculty of the

COMMITTEE ON BUSINESS ADMINISTRATION

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY
WITH A MAJOR IN MANAGEMENT

In the Graduate College

THE UNIVERSITY OF ARIZONA

2000

UMI Number: 9992100

Bell & Howell Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

THE UNIVERSITY OF ARIZONA ®
GRADUATE COLLEGE

As members of the Final Examination Committee, we certify that we have read the dissertation prepared by Gary Scott Mahon

entitled The Performance of Highly Active Problem Solving Strategies in Novel Task Environments

and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

Terence Connolly	12/1/00 Date
Lisa D. Ordòñez	12/1/00 Date
Kenneth W. Koput	12/1/00 Date
	Date
	Date

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copy of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Dissertation Director Terence Connolly	12/5/00. Date

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgement of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: Gary Scott Mahan

ACKNOWLEDGEMENTS

I would like to begin by thanking the members of my committee who helped guide me through this process and provided encouragement along the way. Terry Connolly, Lisa Ordóñez, and Ken Koput were there to assist me whenever I asked. I am grateful for the time and energy that Lisa and Ken put into reading numerous drafts and offering helpful suggestions. My advisor and mentor, Terry Connolly, endured my ramblings at our weekly meetings and was always there to keep me focused. He has an uncanny ability to maintain focus on the details and the big picture simultaneously, without losing sight of the importance of each. His guidance was invaluable.

I also want to thank my family, my wife Brenda and our boy Tyler for their encouragement and support. Finally, to my Mom who passed away in September, thanks. I couldn't have done it without them.

TABLE OF CONTENTS

TABLE OF CONTENTS - CONTINUED

ABSTRACT

A series of simulation experiments evaluated the performance of seven different rule-based problem-solving strategies. Each of the strategies was based on a small set of decision rules that used performance feedback from prior actions to determine future behavior. Several environmental task factors were studied including feedback error, complexity, and system dynamics. The results showed that different strategies perform well in some environments but not in others. No one strategy performed well across the range of environments studied here. The implications for human decision-makers are that in order to be successful across a variety of tasks, a person must possess a broad repertoire of problem solving strategies and know when and how to apply them.

In addition, two laboratory experiments were conducted with human subjects using the same task factors as in the simulation experiments. The findings lend support to a new theory on problem solving in novel task environments. In stable, positive, and non-declining environments, human decision-makers employed a two-stage approach to maximizing their payoff. Behavior in the first stage was characterized by bold actions that were used to explore the environment and gain a basic understanding of the payoff distribution. Approximately one third of the way through the task, subjects changed their problem solving strategy to a more systematic, small step approach similar to the way many of the rule-based simulated subjects behaved. Another interesting result was the inability of subjects to improve their performance in their second run. Relatively minor changes to the task from one run to the next were enough to block the ability to transfer knowledge from the first run to the second. Additionally, 12% of the runs in the

laboratory experiment performed at a level that was less than or equal to what could have been achieved simply by choosing settings at random. These results suggest that subjects performing at this level could have saved a considerable amount of cognitive effort by taking random actions. Additional research is needed to evaluate new task factors, alternative problem-solving strategies, and gain a better understanding of the two-stage approach.

CHAPTER ONE

Introduction

Today's business world is characterized by a fast-changing environment and a high degree of uncertainty about the future, resulting in a constantly shifting paradigm for success. For example, the information technology revolution of the late 20th century has spawned a period of instability and uncertainty in the business and technological environments, not unlike the industrial revolution in the early part of the century. Problem-solving strategies that may have been successful under stable and predictable economic conditions are no longer as effective in a dynamic and fast moving world. To achieve desired performance in a turbulent environment requires a new way of thinking about a firm's product offering, pricing, promotion, production methods, and the way in which the product is delivered to the customer. Churchman (1971) referred to this dilemma as the manager's paradox. On one hand, the manager is under pressure to take quick and decisive action to capitalize on transient opportunities. On the other hand, the manager faces problems where the linkage between actions and outcomes is unknown, implying that action must be taken in order to determine whether the action taken was correct. This raises some interesting questions. How well do people learn and adapt their problem solving strategies to changing environments? Does successful performance in one environment lead to success in other environments? The answers to these questions require new insight into the way that people deal with complexity, instability, and uncertainty in the environment.

A decision-maker confronting an uncertain and complex task has a variety of problem solving strategies at his or her disposal. One way to deal with a multitude of alternative solutions is through exhaustive experimentation and evaluation of each and every possible combination of variables. However, in a practical sense, it is impossible to test all of the available alternatives within a reasonable period of time. Shifting consumer preferences and impatient investors require business executives to find optimal or near-optimal solutions rather quickly. In fast-changing and ill-understood environments, decision-makers cannot base their decisions on a prior strategic understanding of the situation. Instead, they may be forced to adopt what we refer to as highly active strategies (HAS). HAS are a group of strategies that rely on varying degrees of adaptive feedback, systematic experimentation, and trial and error. A number of different HAS have been identified and are discussed below. Though such strategies are generally considered to be desperate, last-resort approaches for decision makers with no strategic understanding of their environments, they can in fact be quite successful, even surpassing analytically optimal approaches under a variety of real-world constraints on information and cognitive capacity (see Lindblom, 1979; Connolly, 1988; Connolly and Wagner, 1988; Quinn, 1989; and Connolly, 1999). The broad purpose of the research proposed here is to explore the performance of different HAS in simulated task environments and to compare this with the performance of real human subjects in similar environments.

Literature Review

This section is organized as follows. First is an extensive review of the prior research in the field of dynamic decision-making. Second is a relatively brief overview of the key findings in the areas of strategic management and organization theory that pertain to the environment and the affect on firm performance. Finally, the work of Gigerenzer and colleagues under the heading of "fast and frugal heuristics" will be summarized.

Dynamic Decision Making

One body of work that is closely related to the focus of this project is the research on dynamic decision-making (DDM). The DDM literature is divided into several subsections. First, the key features and distinguishing characteristics of dynamic decision-making will be discussed to provide a framework. Next, is an overview of the empirical and analytical work in DDM from its beginning in the early 1960's up to the present day. The prior research in dynamic decision-making can be classified, somewhat arbitrarily, into three camps based on the research paradigm. One of these groups includes the work emanating from the University of Chicago and its alumni, noted for their use of computer-based heuristic decision strategies in simulated environments and subsequent performance comparisons with human subjects. The second group of researchers is based primarily in Europe and notable for their use of realistic decision scenarios called Microworlds. Their emphasis is on the effects of task factors and system factors on players' performance. Their simulations are unique in that many of them are run in real-time. The third DDM research stream is centered at MIT, where the focus is on people's ability to learn from system feedback and the construction of mental models.

Edwards (1962) was first to formalize the definition of dynamic decision-making and develop a taxonomy of dynamic decision tasks. Dynamic decision-making as defined by Edwards (1962) is characterized by four conditions. First, dynamic tasks involve a sequence of inter-related decisions made over a period of time. This can be contrasted with static decision-making tasks that typically require only a single, discrete choice. Second, the outcomes of earlier decisions are capable both of generating payoffs and of contributing to the knowledge of the environment. The third feature of dynamic decision tasks is that the information, i.e. feedback, resulting from prior decisions may or may not be relevant to future decisions, depending on the nature of the task environment. Lastly, the environment in which the decision takes place may change over time either as a function of previous decisions (endogenous) or independently of them (exogenous), or both. Edwards (1962) goes on to describe six distinct types of dynamic decision tasks, differentiated by three dimensions. The first dimension is the stability of the environment, which is either stationary or non-stationary depending on whether there are exogenous environmental variables. The second dimension is the extent to which previous choices can affect the task environment, in other words, whether endogenous variables are present in the environment. The third and final task dimension is defined as the relevance of feedback from previous decisions. Because some conditions preclude the existence of others, there are six possible combinations of these three factors. Some examples of empirical work stimulated by Edwards' work are discussed in the remainder of this section.

Hogarth and Makridakis (1981) conducted an experiment to compare the performance of relatively simple heuristic decision strategies to human subjects' strategies in a complex, dynamic task. The experiment revolved around a marketing strategy game covering multiple time periods and a number of different industries. Subjects were assigned to teams and given the task of managing production levels, R&D expenditures, advertising budgets, product pricing, and distribution channels. Unknown to the subjects, the experimenters controlled one team within each industry. Each team's performance was affected by what they did, what their competitors did, and the general economy. The student teams were able to observe prior-period results and make modifications to their strategy, while the experimenters' teams applied either consistent arbitrary decision rules (e.g. estimated sales = prior period sales + GNP growth rate, advertising = 10% of estimated sales) or random arbitrary decision rules (e.g. advertising = 10% of estimated sales +/- some random number). In both conditions, rules were fixed before the start of the game and were not changed regardless of how the teams performed during the game. The results revealed that, on average, the experimenters' arbitrary-consistent teams finished third out of five teams in each industry and outperformed 41% of the student teams. Further, the arbitrary-random teams finished better than 22% of the student teams. The results suggest that for many of the student run teams, the additional time and effort required to plan their strategy for each period did not lead to better performance.

Kleinmuntz and Kleinmuntz (1981) created a medical decision-making task where simulated subjects were presented with a patient complaining of three unknown symptoms and suffering from one unknown disease. Probabilistic relationships were specified, linking the 30 symptoms, five diseases, and 12 treatments. The patient's health was expressed on a scale from 0=dead to 100=cured, with 60 time periods available to improve the patient's health. The patient's health is a function of the disease (a linear downward trend with variability), the treatment(s) applied, and the cost to the patient's health of the diagnostic tests. In each time period, the "doctor" can order a test to check for the presence of a symptom (with immediate feedback) and apply a treatment. Treatment effectiveness depends on both the treatment selected and the patient's disease. There were three computer simulated decision strategies. The first strategy was based on Bayesian revision. The second was a heuristic strategy where hypotheses were developed about the patient's disease and attempts were made to confirm it. Diagnostic tests were chosen based on those most likely to reveal symptoms closely associated with the disease. Treatments were chosen randomly from those that met the minimum criteria for expected effectiveness for that disease. If a treatment was effective, it was retained and used in later periods. If a treatment didn't improve the patient's health it was not used again. The third strategy was termed "generate and test", which was essentially a trial and error approach. Here, treatments were generated at random and tested for favorable results. Treatments were re-applied as long as they were effective. Symptoms were ignored, since there was no attempt to diagnose the disease. Each of the three strategies was simulated for 100 cases. The results show that the Bayesian strategy provided only a marginal improvement over the heuristic strategy, though both substantially outperformed the "generate and test" strategy on four of the five diseases. On one disease, the trial and error strategy outperformed the heuristic strategy. In this particular case, the random strategy produced much greater variation in performance, resulting in more people being killed or completely cured than the heuristic strategy. Overall, the Bayesian strategy required much more computational effort, but produced only a small increase in performance relative to the heuristic strategy. The performance of the random strategy on one of the diseases points out that, as the features of a task change, random strategies are better able to adapt to a new environment. A Bayesian strategy would be unable to perform well under these conditions without a complete re-specification of the task parameters.

Kleinmuntz (1985) used a variation of the medical decision-making task in which it was possible for treatments to have fatal side effects and diagnostic tests and treatments could not be done simultaneously. Two baseline strategies were employed - Bayesian revision and a random treatment strategy. In between the two baseline strategies were various heuristic strategies differentiated by their computational complexity, base-rate utilization, and information acquisition strategy. In one experiment, the heuristic strategies outperformed the Bayesian baseline under certain task conditions, although not significantly. Additionally, it was observed that more complex heuristics did no better than simpler ones. In another experiment, they found that in some cases the random strategy outperformed the heuristic strategy. Overall, the largest performance difference between heuristics was due to information acquisition strategy rather than information processing complexity. This suggests that the ability to select the right variables is more important than the ability to combine them optimally.

Kleinmuntz and Thomas (1987) looked at the relationship between the structure of the decision environment and the use of different decision strategies in an effort to determine when decision-makers would choose action-oriented (treatment) strategies and when they would select judgment-oriented (diagnostic) strategies. A variation of the medical decision task was used. Two treatment factors were manipulated – the presence or absence of a Bayesian decision aid and the possibility that a treatment may have fatal consequences. It was expected that the presence of a Bayesian decision aid would improve the interpretation of outcome feedback, making judgment-oriented (diagnostic) strategies more likely. Further, it was expected that low levels of treatment risk would increase the use of action-oriented (treatment) strategies. There were four dependent variables, including the proportion of patients cured, number of diagnostic tests run, number of treatments, and the elapsed time of each trial in seconds. The results of the within-subjects analysis show significant decreases across trials in the number of diagnostic tests run, number of treatments applied, and elapsed time, suggesting a learning effect. The results of the between-subjects analysis showed significant main effects for both the treatment risk factor and the presence of a Bayesian decision aid, but no interaction. It is also worth noting that subjects with the decision aid took 20% longer to treat each patient. Overall, subjects performed somewhere in-between the random and Bayesian benchmark strategies. However, 92% of the subjects in the low risk condition performed only as well as the random no-wait benchmark. In the random no-wait heuristic strategy, treatments were applied one after another until all treatments were exhausted and no attempt was made to diagnose the disease.

One of the primary European DDM researchers, Brehmer (1990) created a dynamic task called FIRE where subjects must make decisions in real time. To be successful in this task requires that both the correct decisions be made and that they are made at the proper time. The goal of one set of experiments was to determine under what conditions subjects would use a feedback strategy or a feed-forward strategy. A feed-forward strategy is defined as one where the choice of actions is based on predictions of the future state of the system, requiring a more complex mental model than a feedback strategy. In the fire-fighting task, subjects are placed in the role of fire chief who controls the actions of multiple fire fighting units (FFU's). Subjects are shown a computer-generated map of the forest depicting the location of the fire, areas that are burned and unburned, and the position of the FFU's. The fire chief can order FFU's to proceed to specific locations under centralized control or decentralized control. When control is centralized, an FFU cannot fight a fire until it reaches its destination, whereas a decentralized FFU can fight any fire it encounters along the way to the assigned destination. The two objectives of the task are to extinguish the fires and to keep the fire from reaching the fire chief's base. The game ends when all the fires are out or when the fire chief's base burns down. In a series of experiments, feedback delay was manipulated. One group had a 20 second delay, another had a 40 second delay, and the third group had no delay. Overall, subjects' performance improved with experience and they soon learned the best strategy. Over a number of trials, the area lost to fire decreased for all conditions, but performance improvement was more pronounced in the no-delay condition. Subjects in the delay conditions did not try to improve their performance by decentralizing control to the FFU's as was expected. The results suggest that subjects ignored feedback delays that were due to delayed reporting.

Funke (1995) performed an extensive review of the European literature in DDM and classified the research into three areas of study – person factors, task factors, and system factors. Each of these three factors has been shown to affect performance in complex problem solving tasks in different ways. (Note: many of the articles cited in this subsection were written in non-English languages and were interpreted by Joachim Funke).

Person factors are defined as the competency that a subject brings to the task or acquires during the task. Strohschneider (1990, 1991b) compared performance on MORO (peacecorps worker simulation) and VEKTOR (an abstract system) to scores on the BIS (Berlin Intelligence Structure) and found that test intelligence was a predictor of complex problem solving performance. However, Süß, Oberauer, & Kersting (1993) found no significant correlation between any of the BIS scales and performance using the TAILORSHOP simulation. Reither (1981) compared the performance of experienced technical advisors with graduate students using a third world simulation game (DAGU) and found no between-group performance differences, although both groups performed poorly on the system. Interestingly, it was found that the experts used a broader range of actions, were more willing to make decisions from the start, but were less able to adapt to changing task conditions, a condition termed "blindness of the specialists". Vollmeyer and Holyoak (1993) analyzed the strategies of subjects in the BIOLOGY LAB simulation. Subjects' behavior was categorized as using a scientific strategy, a systematic variation of a strategy, or an unsystematic variation of a strategy. When exploring the system or predicting the future behavior of the system, the unsystematic strategy was outperformed by the other two. However, when controlling the system, there were no differences between the three strategies. These findings suggest that different types of knowledge are necessary for different tasks.

Situation factors encompass the context in which the simulation is embedded, independent of the scenario used. In one experiment, Funke and Miller (1988) divided subjects into two groups. Some were told to actively explore the SINUS system, while others passively observed another person using the system. Later, all subjects were asked to control the system. Active intervention in the first session led to better performance in the second session but with lower ability to verbalize the knowledge compared to subjects who only observed in the first session. Köller, Dauenheimer, and Strauß (1993) conducted an experiment where, in the first session, all subjects worked individually on the FUEL OIL DISTRIBUTION task. Subjects were then classified as good or bad performers. In the second session, subjects worked on a similar task called TEXTILE SHOP either individually or in pairs of two good performers or two bad performers. Individuals performed worse than the dyads, although there was no difference in performance between the dyad types. Leuter (1988), using a variation of TAILORSHOP, found no performance difference between individuals and 3-person groups, although individuals seemed to acquire more knowledge about the system than did the groups.

Putz-Osterloh (1981) and Putz-Osterloh & Lüer (1981) looked at transparency using the TAILORSHOP scenario. In the transparent condition, subjects were given a graphical representation showing the inter-relationships among nearly all of the system variables and those in the nontransparent condition were given nothing. Not surprisingly, subjects in the transparent condition outperformed those in the nontransparent condition. Hüber (1994) used the GAS ABSORBER scenario to test the effects of presentation format. Half of the subjects were shown the state of the system in analog format, while the others were shown the system state in digital format. For the dependent variable "quality of system control", subjects using the analog format outperformed those using the digital format.

System factors are the specific attributes of the system that are content-related. One system factor is termed Eigendynamik, which exists when a system changes its state over time as a result of exogenous variables, i.e. independent of any operator input. Funke (1993) used the SINUS simulation and placed subjects in 3 conditions. The control group had no Eigendynamik, while the other 2 groups had varying degrees of Eigendynamik. Results show that increased Eigendynamik leads to a decrease in the dependent variable "quality of system control". Another dependent variable "quality of system identification" was unaffected by the manipulation. The findings suggest that the two dependent variables may tap different cognitive processes. Heineken, et al (1992) tested the effects of feedback delay using the simulation game TEMPERATURE. Feedback delay was varied between none and substantial. Half the subjects were told in advance that there would be feedback delay. They found that the DV "quality of system control" decreased with increasing delay. Prior notification about the delay did not have any effect. After a relatively long period of playing time, even subjects in the longest delay condition were able to control the system. This suggests that feedback delay slows the rate of learning but does not block a person's ability to master the system. Hesse (1982) compared two different cover stories for the same underlying system, EPIDEMIC. In one condition, subjects were told to care for people with a fairly harmless flu, while the others were told to manage a dangerous outbreak of life-threatening smallpox. In the smallpox condition, subjects appeared to be much more involved in the task and took a longer period of time to reach their decisions, although there was no difference in overall performance between groups.

Interactions between person, task, and system variables have also been studied. Rost and Strauß (1993) studied the interaction between presentation format (numerical vs. graphical) and the type of induced mental model (the way that knowledge about the system is internally represented; propositional vs. analog). A short training session was held before the experiment to get subjects thinking in terms of propositions (if-then statements) or in terms of a graphical network (nodes represent variables and lines connecting the nodes represent causal links). It was assumed that the propositional model would best fit the numerical presentation format and the analog model would best fit the graphical format. Using the simulation SHOP, they found that the analog training produced a large performance difference between the presentation formats while the propositional training produced a much smaller difference, providing strong evidence of an interaction between person and system variables.

The DDM research at MIT is led by John Sterman. Sterman (1987) conducted an experiment utilizing a dynamic capital investment task where subjects were placed in the role of manager for the capital-producing sector of the economy. Orders for capital goods were placed at the beginning of each 2 year cycle. In the task, demand is generated from both the goods sector and the capital sector itself, i.e. capital is needed to create capital. Demand from the goods sector is exogenous, while the subject determines the capital sector capital requirements based on desired capacity. The results show that a vast majority of players generate significant oscillations that are far greater than the fluctuations in demand. Only 8% of subjects were able to reach equilibrium over the 70-year period, even though the system reached equilibrium in period five, following the shock in periods three and four. The findings suggest that people may not behave optimally despite perfect knowledge of the system and perfect information. For example, players made basic errors such as ignoring the amount of capital on order, failing to anticipate the time lag in acquiring capital, and failing to anticipate the increase in demand caused by their own orders.

The Beer Distribution Game (Sterman, 1989c) is similar to the lagged capital goods game described above. In this task, subjects try to minimize costs as they manage the production and distribution of this vital commodity. Subjects, including experienced managers, generated costly fluctuations even when consumer demand was constant. Costs, on average, were more than ten times greater than optimal. Most people did not account for the supply line of orders that had been placed but not yet received, causing them to overcompensate for inventory shortfalls. Subjects often blamed their difficulty on exogenous events. When asked to sketch the pattern of customer demand, most drew a large amplitude fluctuation similar to the oscillation they generated.

Diehl and Sterman (1993) examined the performance of subjects in an inventory management task. Time delays and feedback effects were varied from trial to trial. Overall, costs were more than four times greater than optimal, despite financial incentives, training, and repeated play. In the easy condition (no time delay or feedback effects), subjects outperformed the do-nothing rule, but in the difficult condition many under-performed the do-nothing rule. Regression models of subjects' decision rules showed little evidence that subjects adapted their decision rules as the complexity of the task changed. When the environment was complex, subjects seemed to revert to simple rules that ignored the time delays and feedbacks. There was no significant difference in the time taken to make decisions across the different complexity levels, even though the number of variables to consider was much greater in the difficult conditions.

Paich and Sterman (1993) showed that learning in complex dynamic tasks is often poor. A corporate strategy simulation game was developed to simulate the market for a consumer durable product that the subjects manage through the full product life cycle, from launch through decline. The simulation included realistic features such as price elasticity effects, original and replacement demand, competition, learning curves, and capacity acquisition delays. Subjects made price and capacity expansion decisions each quarter for ten simulated years. The subjects played five such trials, each with different characteristics of the market and product. The treatments varied the strength of the feedback loops. Subject performance was compared to a simple decision rule embodying a naive strategy that was insensitive to the feedback structure of the market and the behavior of the competitor. The naive strategy outperforms nearly 90% of the subjects. Performance dropped significantly as feedback complexity increased.

The prior research in the field of dynamic decision-making has produced some interesting results. Hogarth and Makridakis (1981) found that random strategies outperformed 22% of actively managed teams in a business strategy game, suggesting that the time and effort required to plan and execute a strategy did not lead to better performance. Paich and Sterman (1993) observed similar results in an experiment where the random strategy outperformed nearly 90% of the human subjects. Kleinmuntz and Kleinmuntz (1981) observed that a Bayesian optimal strategy provided only a marginal improvement over a heuristic strategy in a medical decision task. In a later experiment using a similar task, Kleinmuntz (1985) found that under certain conditions, a random strategy outperformed the heuristic strategy. In the fire-fighting task, Brehmer (1990) found that subjects' performance improved with experience in the task. System dynamics has been shown to lead to degraded performance (Funke, 1993; Sterman, 1987). Despite perfect knowledge of a system, subjects tended to overcompensate for fluctuations in the environment.

Strategic Management and Organization Theory

The fields of strategic management and organization theory (OT) provide some insight into the relationship between strategy, the environment, and performance at the organization level. This subsection will begin with a brief overview of the five dominant theories in OT and their implications for managerial decision-making under conditions of uncertainty, dynamism, and complexity. It will conclude with a review of some of the key papers in the strategy and organization theory literature.

The five major streams of research in organization theory are resource dependence, transaction cost economics, organizational ecology, institutional theory, and network approaches. Resource dependence (Pfeffer & Salancik, 1978) is based on the assumption that organizations are dependent upon other organizations for critical resources such as capital, workers, raw materials, machinery, etc. Organizations that are resource providers can leverage their power and make demands on those organizations that require the resource. From a decision-making perspective, organizations in need of the resource must decide how to act. It is assumed that organizations desire autonomy and will behave in ways that help them reduce their dependence. In environments characterized by uncertainty and dynamism, the decisions are more complicated. For example, the uncertain availability of capital funding may prompt organizations to utilize alternative sources of funds at higher rates of interest. Rather than analyzing a number of funding options and making an informed choice based on a set of attributes, organizations may be forced to take the first available source of funds in order to avoid a situation where no capital can be found. Transaction cost economics (Williamson, 1975, 1981) posits that firms will attempt to acquire resources using cost as the primary criteria for deciding whether to produce the product/service themselves or purchase the product/service from other organizations. The theory predicts that firms will choose the alternative with the lowest total cost. When the transaction costs of the market are lower than administrative costs, then the choice will be to outsource. In fast-changing, dynamic environments it is possible that organizations will decide to utilize the market for more products and services than when the environment is stable. Although the short-run costs may be higher, outsourcing would enable firms to retain flexibility and reduce uncertainty in turbulent periods. The theories of organization ecology (Hannan & Freeman, 1977, 1989; Carroll, 1984) contend that environmental factors, including dynamism and complexity, determine organizational survival. Although managerial decision-making is constrained, an organization may choose to exploit market niches in order to adapt to a fast changing environment. In dynamic conditions, organizations are expected to increase their spending on research and development for new products to guard against technical obsolescence and enter new market niches. Institutional theory (Powell & DiMaggio, 1991) is based on the premise that key organizations act as institutions within industries or sectors. These institutions can take the form of trade groups, technical consortiums, government agencies, or a lead organization within the industry. Institutions award legitimacy to organizations, thereby enabling access to critical resources. When institutional standards are normative and mimetic rather than mandated, organizations have some discretion over whether these standards are adopted. It is possible that in uncertain environments, organizations may be less likely to choose to adopt the standards established by the institution, compared to environments where there is a low level of uncertainty. This is expected to hold true when the reasons for the uncertainty are attributed to the ineffectiveness of the institution. Network approaches (Davis & Powell, 1992) view organizations as embedded in an interconnected web of relationships. Networks are often built on social rather than economic relationships and serve to enable organizations to mobilize resources. Organizations can choose to belong to the network or not. When the environment is characterized by uncertainty, as is the case when technical standards are evolving, it is expected that firms will be more likely to become part of the network affiliation. In summary, the major streams of research in OT provide some insight into the important environmental factors and managerial decisions that are deserving of future study. Most of the prior research has centered on environments where actions result in changes to the environment and organizations have knowledge of their competitor's actions. This dissertation will explore environments where relatively little is known. Specifically, the experiments will evaluate environments in which actions do not change the state of the environment and the actions of competitors are unknown.

Contingency theory (Thompson, 1967; Lawrence & Lorsch, 1967) suggests that successful performance depends on finding the proper alignment between the environment and a firm's business-level strategy (Miller, 1987b) and its organization structure (Venkatraman & Prescott, 1990; Priem, 1994). This implies that the highest performing firms are those that are the best at matching their strategy and structure to the features of the environment. However, two important considerations have limited the analytical research on strategy-environment co-alignment. First are the problems surrounding the conceptualization and operationalization of environments and strategy. Second is the problem of finding an appropriate analytical technique for systematically measuring the level of co-alignment and its impact on performance (Venkatraman & Prescott, 1990).

The remainder of this subsection will briefly review the research on the incremental approach to strategy formulation and implementation as well as take a look at two key environmental factors, dynamics and complexity. Dess and Beard (1984) identified munificence (defined as the capacity of the environment to support sustained growth), complexity, and dynamism as the three most important environmental factors affecting firm performance.

Disjointed incrementalism (or, less formally, "muddling through") is the name given to the process used by public policy makers in their attempt to find solutions to complex problems. "Complex problems cannot be completely analyzed and we therefore require strategies for skillful incompleteness..." (Lindblom, 1979). Disjointed incrementalism is based on a sequence of trial and error experimentation, where analysis is necessarily limited to a relatively small number of policy alternatives.

Campbell (1969) advocates a quasi-experimental approach to matters of pubic policy and social reform. In attempting to solve social problems, administrators are often confronted with ambiguity, both in terms of uncertain action-outcome linkages and in the evaluation of outcomes. The term "staged innovation" is used to describe the deliberate spreading out of social reform programs among limited segments of the affected population or by geographic area. The intent of staged innovation is to learn whether the specific programs are effective before they are deployed on a widespread basis. Those solutions that are successful are either retained as is or modified somewhat and those that do not perform as intended are discarded.

In a similar vein, Quinn (1989) has studied incremental approaches to strategy in a business context. As expected, managers were found to utilize very rigorous strategic planning processes, though they were also apt to let strategies evolve incrementally over time. Quinn (1989) provided several reasons were given as to why executives formulate strategy in this way. Among them is the ability to remain flexible as conditions change, enabling the business to opportunistically pursue multiple goals. Several U.S. Department of Defense contractors, including IBM and Xerox, have developed "phased program planning systems" where final decisions are made only on the current phase of a project, while future options remain open to subsequent evaluation. Major commitments of resources, such as manufacturing plants and capital equipment, are deferred for as long as possible, so that all available and relevant information may be taken into account before the final decision is made. Overall, the line separating strategy formulation from implementation becomes blurred. In many organizations, these two activities often occur nearly simultaneously.

The so-called environmental view of strategy began in the mid 1960's (Emery and Trist, 1965; Lawrence & Lorsch, 1967; Terreberry, 1968, Perrow, 1970). Simply stated, if organizations are to be successful, they must learn to adapt to their environment. One environmental factor that strongly influences organizational performance is complexity. Child (1972) defined complexity as "the heterogeneity of and range of an organization's activities". Duncan (1972) further refined the complexity dimension of the environment as "the number of factors taken into consideration in decision-making". For example, in a simple environment, the decision factors are relatively few in number and similar to one another. Duncan (1972) uses the example of a production scheduler who manages a handful of part numbers and interacts with the marketing and manufacturing departments. On the other hand, in a complex environment, the decision-maker must take into account a wide variety of factors, both internal and external, when making a decision. An example would be members of a planning department who take into account a very large number of considerations when making a decision and who must interact with a multitude of organizational entities. Dess and Beard (1984), building on the work of Child (1972), defined complexity as encompassing two subdimensions: heterogeneity and dispersion. The homogeneity-heterogeneity subdimension is a measure of the range of an organization's activities. A single product firm in a mature industry exists in a relatively homogenous environment, while a multi-national firm engaged in a variety of industries with a wide range of customer needs can be thought of as operating in a heterogeneous environment. The concentration-dispersion subdimension is an indicator of the total number of organizations that exist within a firm's competitive environment. For example, the industry encompassing commercial aircraft manufacturers would rank very low on this dimension, while the cotton textile industry would be positioned relatively high on the scale. Aldrich (1979) argued that an increase in organizational density increases the need for strategic activities. Starbuck (1976) contended that as density increases, organizational interdependence would also increase.

Duncan (1972) defined the dynamic dimension of the organizational environment as "the degree to which the factors of the decision unit's internal and external environment remain basically the same over time or are in a continual process of change". This includes such factors as personnel, structure, customers, suppliers, competitors, technology, and government entities. In a static environment these factors change very little, while at the other end of the continuum, in a dynamic environment nearly everything is in a constant state of flux. Perrow (1972) defined a dynamic environment as a "turbulent, ever-changing environment". Aldrich (1979) divided environmental dynamics into two subdimensions, stability-instability and turbulence. Stability-instability refers to the degree of turnover in environmental elements. Aldrich (1979), however, noted that high levels of change could be patterned and thus predictable. The second dimension, turbulence, refers to an increasing rate of change of interconnections between elements and trends in the environment. In sum, it is not only the size of the change that creates a dynamic environment; it is also the rate of change.

Miles, Snow, and Pfeffer (1974) point out that it is not only the rate of change that matters; it is also the predictability of the changes. Thus, dynamism refers to change that is hard to predict and heightens the uncertainty level of the organization's members. Nelson and Winter (1982) argued that a dynamic environment exists when changes take place in external market conditions or occur as a result of innovation within the industry itself. Eisenhardt and Bourgeois (1988) researched the microcomputer industry in mid to late 1980's and found that dynamic conditions exist in high velocity environments with "rapid and discontinuous change in demand, competitors, technology, or regulation, so that information is often inaccurate, unavailable, or obsolete". Eisenhardt (1989) observed that dynamic environments are fast-paced settings with a "tumult of technical change", requiring managers to be fast decision makers to keep pace with the changes and perform better. Reger & Palmer (1996) studied competitive practices in dynamic environments and noted that managers view competition quite differently in turbulent environments.

The key findings of the strategy literature are summarized below. Venkatraman & Prescott (1990) and Priem (1994) found that the highest performing firms are those that are the best at matching their strategy to the features of the environment. Child (1972) stated that if organizations are to be successful, they must learn to adapt to their environment. Incremental approaches to strategy (Lindblom, 1979; Quinn, 1989) are based on the observation that many problems are too complex to be completely analyzed. Strategies are based on trial and error experimentation and evolve incrementally over time in contrast to a fully developed top-down approach. This enables organizations to remain flexible and better respond to changing conditions.

Fast and Frugal Heuristics

Fast and frugal heuristics (Gigerenzer & Goldstein, 1996; Gigerenzer & Todd, 1999) are decision rules or processes that require a minimum of time, knowledge, and computation to make choices in real environments. Surprisingly, such heuristics can be as accurate as strategies that use all available information and extensive computations. They can be used to solve problems of sequential search through options or to make choices between simultaneously available options. Members of the family of fast and frugal heuristics all rely on three basic steps: rules for guiding the search process, rules for knowing when to stop the search process, and rules for making a choice.

Decisions are made between alternatives based on information about those alternatives. In many situations, the alternatives and pieces of information will need to be found through active search. For instance, cues can be searched for in a random manner, or in order of some pre-computed criterion related to their usefulness or based on a recollection about which cues worked previously when making the same decision. Similarly, the search for alternatives can be random or ordered.

The search for alternatives or information must be terminated at some (preferably early) point. The method for determining when to stop search should not be overly complicated. For example, one simple stopping rule is to cease searching for information as soon as the first cue or reason that favors one alternative is found. For searching through alternatives (rather than cues), simple aspiration-level stopping rules can be used, as in Simon's principle of satisficing (Simon, 1956b, 1990).

Once the appropriate alternatives or information have been found, a final set of heuristic rules is used to make the decision based on the results of the search. For instance, a decision or inference can be based on only one cue or reason, whatever the total number of cues found during search. Such one-reason decision-making does not need to weight or combine cues, and so no common currency between cues needs to be determined. Decisions can also be made through a simple elimination process, in which alternatives are thrown out by successive cues until only one final choice remains.

Fast and frugal heuristics can be divided into four categories: ignorance-based decision-making, one reason decision-making, elimination heuristics, and satisficing heuristics. Ignorance-based decision-making uses the recognition heuristic to select one option from two possibilities, according to some criterion on which the two can be compared. When choosing between two objects, if one is recognized and the other is not, then select the former. Employing the recognition heuristic can lead to the surprising less-is-more effect, in which an intermediate amount of knowledge about a set of objects can yield the highest proportion of correct answers. In one experiment, subjects were asked to choose the larger of two German cities in terms of total population. Surprisingly, those who knew more about the cities (German citizens) actually performed worse than those who knew less (Americans) (Gigerenzer & Goldstein, 1996).

One-reason decision-making uses a stopping rule that terminates the search for information as soon as enough has been gathered to make a decision. The rule says "stop looking for cues as soon as one is found that differentiates between the two options being considered". (1) Select a cue dimension and look for the corresponding cue values of each option; (2) compare the two options on their values for that cue dimension; (3) if they differ (e.g., if one value is larger or if there is positive information for one option but not for the other), then stop and choose the option with the cue value indicating the greater value on the choice criterion; (4) if they do not differ, then return to step 1 to look for another cue dimension. One example is called the "Take The Best" heuristic, which searches for cues in the order of their validity, i.e. how often the cue has indicated the correct versus incorrect options. "Take The Last" looks for cues in the order determined by their past success in stopping search, so that the cue that was used for the most recent previous decision is checked first when making the next decision. Finally, the Minimalist heuristic selects cues in a random order.

Elimination heuristics are used for multiple-option choices where choices must be made among several alternatives. In situations where each available cue dimension has fewer values than the number of available alternatives, one-reason decision-making will usually not suffice, because a single cue will be unable to distinguish between all of the alternatives. For instance, knowing whether or not each of 15 cities has a river is not enough information to decide which city is most habitable. In these instances, successive cues are used to eliminate more and more alternatives and thereby reduce the set of remaining options, until a single option remains. The QuickEst heuristic is one of the elimination heuristics and is designed to estimate the values of objects along some criterion while using as little information as possible. QuickEst is designed to work well in environments characterized by a J-distribution, where there are many more objects at one end of a criterion range than at the other. To exploit this environmental structure, QuickEst first looks at a cue that separates the most common objects from all of the others (e.g., because most small cities in Germany do not have a professional soccer team, this cue should be one of the first checked when estimating a German city's population). QuickEst then looks at the next cue that separates the remaining most common objects from the rest, and so on until an estimate can be made.

Satisficing heuristics are applicable to situations where alternatives themselves (as opposed to cue values) take time to find, appearing sequentially over an extended period of time. In this type of choice task, a fast and frugal decision-maker must limit information search and have a stopping rule for ending the search for alternatives themselves. An aspiration level is set for the selection criterion being used and the search for alternatives is stopped as soon as the aspiration is met.

There are two primary reasons why fast and frugal heuristics are effective. First, fast and frugal heuristics can benefit from the way information is structured in the environment. The QuickEst heuristic relies on the skewed distributions of many real-world variables such as city population size. Traditional statistical estimation techniques would either ignore or erase this distribution by normalizing the data. Standard statistical models aim to be as general as possible, so they make as broad and as few assumptions as possible about the data to which they will be applied. Second, the heuristics utilize non-compensatory rules. The Take The Best heuristic equals or outperforms a linear decision strategy when information is non-compensatory, i.e. when the potential contribution of each new cue falls off rapidly.

In summary, fast and frugal heuristics are not simply scaled-down versions of optimal strategies. Real world demands often require making accurate decisions in a minimal amount of time, using a minimal amount of information. Fast and frugal heuristics have shown that they can perform as well as and sometimes better than so-called optimal strategies in real-world situations.

There are a number of similarities as well as important differences between highly active decision strategies and the work on fast and frugal heuristics (FFH). They are similar in the sense that they are simple, rule-based problem solving strategies that perform well across a variety of different task environments. There are, however, some significant differences. They differ in the source of the information used to make decisions, the dynamic nature of the task environment, and the number of alternatives available to the decision-maker. In the FFH research, information is retrieved from a subject's memory in order to make the determination as to which of two cities is larger. In contrast, HAS derive their information entirely from outcome feedback. Secondly, thus far the FFH research has dealt only with static decision tasks where the outcome of previous choices has no affect on subsequent decisions. In dynamic decision-making, the decision-maker's knowledge of the environment is constantly being updated through outcome feedback, which in turn, affects future decisions. The third major difference between FFH and HAS is the number of options available to the decision-maker. In a typical FFH task, a subject is asked to choose from among two alternatives, for example selecting the larger of two Germany cities. The task used in this project offers the decision-maker over 3,000 possible alternatives in each time period from which to choose.

Summary of Literature Review

The preceding review provides a good foundation for future research in dynamic decision-making. The prior research sheds some light on a number of promising areas for additional work, including several important task environment factors that are of interest for this project. These task/environment variables include feedback accuracy, task complexity, and system dynamics.

Feedback accuracy, also referred to as feedback error or feedback quality in the literature, is inherent in many DDM tasks but has received comparatively little attention from researchers as an independent variable (Brehmer, 1992). One exception is the work of Connolly and Miklausich (1978) who observed that increasing levels of feedback error led to decreased performance in diagnostic inference tasks. As Brehmer (1992) noted, the presence of even small amounts of feedback error leads to variation in the observability of the system, which in turn may lead to performance degradation in the task.

The importance of complexity as an important variable for further study is clear in both the DDM and the strategy literature. Brehmer (1992) looked at complexity in an advanced version of the FIRE task where subjects were confronted with multiple conflicting goals. In the strategy literature, complexity has been found to be one of the three major environmental factors affecting organization performance. Dynamics and organization structure are the others (Eisenhardt, 1992). Generally, it has been found that the greater the number of variables and/or goals in the environment, the lower the performance.

System dynamics, a.k.a. Eigendynamik (Funke, 1993, 1995), can be thought of as the degree of volatility or instability in the system, independent of the prior actions of the player. Higher levels of system dynamics have led to significant decreases in overall task performance. At the organization level, dynamics is a key determinant of performance (Eisenhardt, 1992). Further, dynamics can be separated into two components, frequency and size. Aldrich (1979) found that both the rate of change and the amount of change affect organization performance.

Research Question

What are the effects of feedback error, complexity, and dynamics on behavior and performance in novel task environments?

Of interest, is the behavior and performance of heuristic decision strategies and human subjects in a variety of task environments. The highly active problem solving strategies serve as simulated subjects and consistently apply preset decision rules to achieve the goal. Human subjects performing similar tasks will provide an interesting contrast to the simulated subjects. Using the data obtained in the laboratory experiments, a behavioral model of human behavior will be developed.

CHAPTER TWO

Task

In the task environment used in this study, a payoff, P_t, is a simple linear function of the decision variable, $X_1 \ldots X_5$, plus trend and error terms. The subject's goal is to maximize the payoff, or profit, over the duration of the task. The task environment is not readily observable by the decision-maker. The actions taken in each period produce both payoffs and information about the task environment. The strategist, whether computerized or human, must rely entirely on outcome feedback for information about the environment.

The task consists of a production function with five variables, each having five intensity levels for a total of 3,125 combinations of possible settings in each time period. The production function also includes an exogenous trend term d_t, which is unaffected by the actions of the decision-maker. In addition, there is an error term e_t, normally distributed with a mean of zero and a standard deviation s_e. In sum, four factors affect the payoff in each time period, 1) the coefficients of the task variables w_i, which are fixed for the duration of the task, 2) the settings applied by the decision-maker X_{jt}, 3) the exogenous trend term d_t, and 4) the error term e_t. The production function takes the form:

$$P_t = w_1X_1 + w_2X_2 + w_3X_3 + w_4X_4 + w_5X_5 + d_t + e_t$$

The length of the task is thirty time periods. The sequence of events is as follows. Initially, the decision-maker is presented with an unknown task environment. In period one, the decision-maker sets the intensity level for each of the five variables. The production function computes the payoff and makes the result known to the decision-maker. In subsequent periods, the decision-maker may choose to change or retain the settings for each of the five variables. This process of taking action and receiving performance feedback continues for the remainder of the task.

Task Factors

Feedback error is manipulated by varying the magnitude of the standard deviation s_e of the error term e_t. Varying levels of s_e will be examined in the computer simulations. The level(s) that produce interesting results in the computer simulation will be used in the lab experiment.

Complexity can be operationalized in at least three different ways. One approach is to manipulate the number of task variables in the production function. For example, a low complexity task would contain a small number of task variables, whereas a moderate to high complexity task would have more variables. A second definition for complexity is the presence of interactions among task variables in the production function. A third way to operationalize complexity is through the use of non-linearities. Simulations using these three representations of complexity have shown that non-linearities have the potential for generating the most interesting results. Therefore, this is the form of complexity that will be evaluated in the computer simulations and lab experiment.

System dynamics is manipulated using the d_t term in the production function. In the computer simulation, the effects of a linear increasing trend and a linear decreasing trend will be examined. Other forms of system dynamics using non-linear functions will be tested in subsequent experiments following completion of the dissertation.

Subjects - Highly Active Strategies

Highly active strategies (HAS) are simple, yet effective problem-solving strategies. Each HAS is based on a set of decision rules that determine how performance feedback is interpreted and the action to be taken in subsequent periods. The strategies differ from one another by the ways in which feedback is used and the rules for determining the next action. At present, we have investigated seven HAS: hillclimber (HC), explorer (EXP), random-best-yet (RBY), a completely random strategy referred to as Forrest Gump (FG), as well as variants of the first three strategies that incorporate multiple period feedback aggregation. Additional highly active problem-solving strategies will be evaluated in subsequent research. The seven highly active decision strategies are differentiated from one another on the following dimensions. Step size is the change in intensity level of a single variable from one period to the next. A period-to-period change in intensity level of one is considered a small step size and a change of two levels or more is defined as large. Search randomness is a measure of the degree to which the next action is systematic or random. A systematic search is one that is determined by the decision rules of the particular HAS and a random search is just that, random. Feedback aggregation is the averaging of outcome feedback over multiple periods. Aggregation can be an effective response to errors in feedback quality. However, there is a tradeoff for improved accuracy. In fixed time period tasks, feedback aggregation increases the accuracy of performance feedback while simultaneously reducing the scope of the search for higher payoffs. See Table 1 for an overview of the seven HAS to be evaluated.

For illustration, the hillclimber and random best yet strategies will be described in greater detail. The hillclimber strategy begins by setting all variables to their lowest intensity level in period one to establish a baseline payoff. In period two, a single variable is selected at random and increased by one intensity level. If the payoff in period two is higher than the payoff in period one, the intensity level of the active variable is increased by one more increment in period three. Provided that the payoffs continue to increase, the intensity level of the active variable is increased by one until it reaches its maximum. Another variable is then selected at random and the process is repeated. If the payoff in the current period is lower than that of the prior period, the intensity level of the active variable is decreased by one increment, a new variable is selected at random, and its intensity is increased by one level. This process continues until all five variables have been evaluated. At this point, the intensity levels for all variables are held constant for the remaining periods in the simulation.

The RBY strategy begins in period one with a random setting of intensity levels for each of the five variables. These settings are then stored in memory as those that generate the best-yet payoff. In period two, another random combination of intensity levels is chosen and the payoffs in the first and second period are compared. If the payoff in period two is higher than the payoff in period one, the period two settings become the new best-yet combination. In the remaining periods of the simulation, the RBY strategy alternates between applying the best yet settings (harvesting) and a random search for a higher payoff (exploring). Each time a new combination of settings produces a higher payoff than the payoff generated by the best-yet settings, it becomes the new best-yet combination.

The seven HAS differ on the amount of memory required and the number of decision rules needed to execute the strategy. Of the four primary strategies, the EXP strategy requires the most memory and decision rules. The EXP strategy includes all of the logic contained in the hillclimber strategy plus additional rules and memory to establish the proper starting points for each variable. The hillclimber strategy requires the second highest number of decision rules and memory. The RBY strategy requires fewer decision rules and memory than the HC strategy, due to the use of random rather than systematic search rules. Forrest Gump is the most frugal strategy, requiring only one rule to choose settings at random and has no need to remember prior payoffs or settings. The feedback aggregation variants of the EXP, HC, and RBY strategies require additional rules and memory than their non-aggregating counterparts. In order from highest to lowest, the decision rules and memory requirements for the seven HAS are as follows: EXP2, EXP, HC2, HC, RBY2, RBY, FG.

Hypotheses

A number of hypotheses have been developed based on the literature as well as my own computer simulation work conducted in developing this proposal.

Because systematic search strategies rely entirely on prior outcome feedback to determine the next action, they are sensitive to errors in the accuracy of the feedback. When the amount of error exceeds the payoff from increasing a variable by a single intensity level, strategies based on small step systematic search rules become confused.

In task environments that contain feedback error, it is not possible to disentangle the payoff attributable to the actions taken from the error. Random search strategies are not susceptible to errors in feedback except at extremely high levels. At low to moderate levels of feedback error, the performance of random search strategies will be unchanged. This leads to the following two hypotheses:

Hypothesis 1a: Systematic search strategies will outperform random search strategies in task environments containing only trace amounts of feedback error.

Hypothesis 1b: The presence of feedback error in the task environment will enable random search strategies to equal or outperform small step systematic search strategies.

The use of non-linear coefficients to represent system complexity should produce no change in the performance of the random search strategies (FG, RBY, and RBY2), but is expected to have a differential effect on the systematic strategies (HC, HC2, EXP, and EXP2). Systematic search strategies adjust the intensity level of a single variable in an attempt to find the setting that produces the maximum payoff for that variable. Both the hillclimber and explorer strategies establish starting points for each variable prior to beginning the small step systematic search for the maximum payoff. In complex environments with non-linear coefficients, systematic strategies are susceptible to getting stuck on local maximum peaks that produce less than optimal payoffs. In cases where the explorer strategy selects the maximum intensity level as its starting point, it is expected that:

Hypothesis 2a: The explorer strategy will outperform the hillclimber strategy when the non-linear function generates declining payoffs as the intensity level increases from the minimum setting to the next higher setting.

Hypothesis 2b: The hillclimber strategy will outperform the explorer strategy when the non-linear function generates declining payoffs as the intensity level decreases from the maximum setting to the next lower setting.

The effects of system dynamics on small step systematic search strategies will be similar to those caused by feedback error. The outcome feedback includes both the payoff attributable to the actions taken as well as the effects of the system dynamics. The strategist is unable to separate the two components, so small step systematic search strategies cannot properly interpret performance feedback. System dynamics are not expected to affect the performance of random search strategies.

Hypothesis 3: Moderate to high levels of system dynamics in the task environment will enable random search strategies to equal the performance of systematic search strategies.

One technique for coping with feedback error is through the use of feedback aggregation. This technique involves taking identical actions in two or more consecutive time periods. The payoffs from each period are averaged together in an attempt to disentangle the payoff due to the action taken from the feedback error. The use of feedback aggregation is expected to enable systematic search strategies to perform better in environments containing feedback error than the same basic strategy that does not aggregate performance feedback.

Hypothesis 4: In task environments containing feedback error, strategies that employ two-period feedback aggregation will outperform the same basic strategy that does not use feedback aggregation.

Results

A series of four simulation experiments were run to assess the relative performance of each highly active strategy (HAS) under varying levels of feedback error, system dynamics, and complexity. In the first experiment, the baseline performance for each strategy was established. Experiments two through four manipulate the level of each task factor and compare the results to the baseline and to other strategies in the same environment. Each of the four experiments produced interesting results.

Experiment One: Baseline

In the baseline condition, all task factors were set to their minimum levels. Each of the seven highly active strategies were run fifty times through the simulated task. The dependent variable is defined as the cumulative payoff at the end of the task expressed as a percent of the potential maximum payoff. This is calculated by taking the cumulative payoff at the end of period thirty and dividing it by the payoff that could have been achieved if each variable were set at its optimum intensity level for every time period in the task.

The results of this experiment are summarized in Table 2. In this and subsequent experiments, the completely random strategy (FG) will serve as the benchmark against which the performance of the other HAS will be compared. In the baseline environment, the FG strategy performed at 38% of maximum ($\pm$ 1.5%, α = .01). In increasing order of performance, the hillclimber strategy with a two period feedback aggregation rule (HC2) achieved 43% of maximum ($\pm$ 0.0%, α = .01). The two random-best-yet strategies (RBY and RBY2) were virtually equivalent to one another, producing payoffs of 51% ($\pm$ 2.3%, α = .01) and 48% ($\pm$ 3.1%, α = .01) respectively. Next best was the explorer strategy with two period feedback aggregation (EXP2), which earned 60% of maximum ($\pm$ 0.5%, α = .01). The hillclimber strategy (HC) performed at 71% ($\pm$ 0.0%, α = .01) and the explorer strategy (EXP) outperformed all other strategies at 77% of maximum ($\pm$ 0.8%, α = .01).

Not surprisingly, the six rule-based decision strategies outperformed the completely random strategy (FG) in the baseline environment. Also, it was expected that the non-aggregation strategies (HC, EXP, and RBY) would equal or outperform their feedback aggregating counterparts (HC2, EXP2, and RBY2). The aggregation strategy variants trade off reduced scope of search for improved feedback accuracy. In environments with accurate feedback, the evidence suggests that the aggregation rule does not improve performance. Both non-aggregating systematic strategies (HC and EXP) outperformed the random-best-yet (RBY) strategy, lending support to Hypothesis 1a that systematic strategies work best in stable, accurate environments.

Experiment Two: Feedback Accuracy

The second experiment examined the performance of the seven highly active strategies in the presence of feedback error. The error term in the production function was normally distributed with a mean of zero and a standard deviation, s_e. Feedback error was evaluated at three different levels: $s_e = 3$, $s_e = 6$, and $s_e = 12$. Each HAS was run fifty times in each task environment.

The results of this experiment are summarized in Table 3. Performance of the three strategies that use random search rules (FG, RBY, and RBY2) were unaffected by high levels of feedback error relative to the baseline environment. By design, random search strategies do not rely on prior period performance to determine the next action(s) to the same extent as the systematic search strategies (EXP, EXP2, HC, and HC2). In fact, the FG strategy totally ignores performance feedback. Although there is no change in the mean performance scores of the random search strategies, there is an increase in the confidence intervals at higher levels of error due to the effects of the error on the distribution of payoffs.

The performance of both hillclimber strategies (HC and HC2) decreases as the magnitude of feedback error increases. The performance drop occurs because the payoff attributable to the action taken and the error are confounded and cannot be differentiated. Technically, this happens when the error term is the opposite sign and larger than the payoff from increasing the active variable by one intensity level. When this condition exists, the hillclimber strategies will set and maintain the intensity level of positive coefficient variables at something less than their maximum value. Similarly, negative coefficient variables will be fixed at something greater than the minimum level. For example, if the active variable has a coefficient of +3 and the standard deviation of the error term is 3, the hillclimber strategies will make incorrect decisions approximately one-sixth of the time. As the size of the error term increases, the percentage of incorrect settings also increases. One of the more interesting findings of this experiment is that at moderate to high levels of feedback error ($s_e = 6$ and $s_e = 12$), the performance of the HC and HC2 strategies actually drops well below that of the completely random strategy (FG). Thus, Hypothesis 1b that predicted the random strategy (FG) would outperform the small-step systematic strategy (HC), is supported.

The explorer (EXP) strategy outperformed all other strategies when error levels were low ($s_e = 3$) to moderate ($s_e = 6$), although performance was lower than in the baseline environment. When error levels are high ($s_e = 12$), performance declines substantially for both EXP and EXP2 (45% and 38%, respectively). Low to moderate levels of error did not affect the large step size moves used to fix the starting positions of the variables in the early periods of the task. However, when the level of feedback error is high, the payoff differential between the maximum and minimum intensity level of a single variable can be less than the error term. When the error is the opposite sign and larger than the payoff differential, the initial starting points are incorrectly set. In later periods when the explorer strategies use small step sizes to fine-tune the settings, they are susceptible to the same potential problems as the hillclimber strategies. At high levels of feedback error, performance of the explorer strategies is only equal to or slightly higher than that of the random search strategies (FG, RBY, and RBY2).

Surprisingly, the feedback aggregating strategy variants (RBY2, HC2, and EXP2) failed to outperform their non-aggregating counterparts under all error conditions. Both the HC2 and RBY2 strategies equaled the performance of their counterparts. However, the EXP strategy outperformed the EXP2 strategy by seven to ten percentage points across all levels of error. Based on this evidence, Hypothesis 4 was not supported. Hypothesis 4 had predicted that strategies employing a feedback aggregation rule (HC2, EXP2, and RBY2) would outperform their non-aggregating counterparts (HC, EXP, and RBY) in environments containing feedback error. In this time-constrained task, the tradeoff between improved feedback accuracy and reduced scope of search was not effective in improving performance.

Experiment Three: Complexity

This experiment evaluated the performance of highly active strategies in complex task environments. To create a complex environment, one of the five variables in the production function was changed from a linear to a non-linear variable. The other four coefficients were unchanged. Two different non-linear functions were used, each containing both a local peak and a global peak. In function #1 the local peak was positioned at the highest intensity level setting for the variable, while in function #2 the local peak was located at the lowest setting. In both functions the global peak was at the middle intensity level. The seven HAS were run fifty times in each environment.

In the environment containing function #1, the performance of the FG, RBY, and RBY2 strategies were unchanged relative to the baseline environment. As expected, these random-search strategies were not affected, either positively or negatively, by the non-linearity.

The performances of the hillclimber strategies in complex environment #1 were up slightly from the baseline performance. HC performed at 74% of maximum (± 0.0%, $\alpha = .01$) versus 71% (± 0.0%, $\alpha = .01$) in the baseline environment. The HC2 strategy was up from a baseline performance of 43% (± 0.0%, $\alpha = .01$) to 49% (± 0.0%, $\alpha = .01$) in complex environment #1. The slight performance increase was the result of the strategy taking fewer time periods to reach the maximum. Because the hillclimber strategies start at the lowest intensity level of each variable, they were unaffected by the local peak on the variable with the non-linear coefficient. In this environment, the HC strategy outperformed the EXP strategy by 10%, providing support for Hypothesis 2b. Hypothesis 2b predicted that when the non-linear function produced declining payoffs as the intensity level decreased from the maximum level to the next lower level.

Both explorer strategies experienced a performance decline. EXP lost 14% compared to the baseline, while EXP2 dropped 9%. The explorer strategies correctly set the starting point at the highest intensity level for the variable with the non-linear coefficient. However, EXP and EXP2 became trapped at the local peak and were unable to reach the global peak, resulting in lower performance. Although the performance of the EXP strategy was down substantially, it was still able to outperform the three random search strategies.

In complex environment #2, once again the performance of the random search strategies was unaffected by the complexity manipulation.

The performance of the HC strategy was down by sixteen percentage points compared to the baseline environment. (55% ± 0.0, $\alpha = .01$ versus 71% ± 0.0, $\alpha = .01$). The HC2 strategy also performed worse than it did in the baseline condition, although it was down by a smaller margin than the non-aggregating hillclimber. The performance decline is attributable to the inability to reach the global peak of the variable with the non-linear coefficient. Instead, both hillclimbers were stuck at the local peak and unable to move.

Both explorer strategies produced virtually identical results compared to their performance in the baseline environment. The EXP strategy outperformed the HC strategy in complexity environment #2 supporting Hypothesis 2a that predicted the hillclimber strategy would outperform the explorer strategy in complex environments with a specific type of non-linear coefficient.

Feedback aggregation was not effective in improving performance in the complex environments evaluated in Experiment Three. Both non-aggregating systematic search strategies (HC and EXP) outperformed their feedback aggregating counterparts (HC2 and EXP2), while the performance of RBY2 was equivalent to that of RBY.

Experiment Four: System Dynamics

The effects of system dynamics on strategy performance were evaluated in this experiment using a linear increasing trend function and a linear decreasing trend function. Each trend function was tested with two different slope values for a total of four dynamic environments. Each highly active strategy was run fifty times in each environment.

Performance of the random search strategies (FG, RBY, and RBY2) was unchanged relative to the baseline environment for all types and levels of system dynamics in this experiment.

In dynamic environments with an increasing trend, hillclimber performance was down slightly when the slope was +2 (65% ± 2.5%, $\alpha = .01$) and down substantially when the slope was +4 (44% ± 0.0%, $\alpha = .01$) relative to performance in the baseline environment (71% ± 0.0%, $\alpha = .01$). When the slope of the trend is equal to or greater than the negative coefficient variable in the production function, the HC strategy has the potential to make incorrect choices. Under these conditions, the hillclimber increased the intensity level of the variable with the negative coefficient. When the slope was +2, the negative coefficient variable was increased half of the time (due to trace levels of error that act as tie-breakers) and when the slope was +4 it was increased all of the time.

The Explorer strategy also experienced a performance decline in environments with an increasing trend. When the slope was +2, performance was 59% of maximum (± 1.6%, $\alpha = .01$) and when the slope was +4, performance dropped to 25% (± 0.0%, α = .01). This compares to a baseline performance of 77% (± 0.8%, $\alpha = .01$). The setting of the starting points in the early periods was unaffected by the increasing slope. However, the EXP strategy was vulnerable to all of the same problems as the HC strategy. In addition, the EXP strategy incorrectly decreased the intensity levels of the positive coefficient variables when the slope of the trend was equal to or greater than the value of the coefficient. When the slope was +2, one positive coefficient variable was decreased all of the time and two variables were decreased half of the time. When the slope was +4, all four of the positive coefficient variables were decreased all of the time.

In dynamic environments with a decreasing trend, the hillclimber strategy performed at 39% of maximum (± 1.3%, $\alpha = .01$) when the slope was –2 and at 9% (± 0.0%, $\alpha = .01$) when the slope was –4. When the slope of the dynamic trend is negative, the hillclimber cannot differentiate between the payoff from increasing a variable by a single intensity level and the effects of the underlying trend. As a result, when the slope was -2, one positive coefficient variable was never increased and two positive coefficient variables were only increased half of the time. When the slope was –4, none of the four positive coefficient variables were increased.

Performance of the explorer strategy in decreasing dynamic environments was only slightly affected by the underlying trend. When the slope was –2, performance was virtually unchanged at 80% (± 0.4%, $\alpha = .01$). When the slope was –4, the EXP strategy performed at 73% (± 0.0%, $\alpha = .01$), a slight drop relative to the baseline. The performance declined as a consequence of the starting point for one of the variables being incorrectly set 50% of the time. The starting point has the potential for being incorrectly set when the payoff differential between the minimum and maximum settings of a variable is the opposite sign and equal to the magnitude of the slope. Based on the evidence presented above, Hypothesis 3 is only partially supported. Hypothesis 3 predicted that moderate to high levels of dynamics will reduce the performance of the systematic strategies (HC and EXP) to that of the random baseline strategy (FG).

Discussion

These experiments produced some interesting, and in some cases, surprising results. One somewhat counterintuitive observation is that under certain conditions it is possible for systematic strategies to perform worse than the random strategy (FG). This was true in environments with imperfect feedback accuracy and in dynamic, fast-changing environments. At moderate to high levels of feedback error, FG clearly outperformed HC. In dynamic environments with a moderately decreasing trend, the completely random strategy (FG) again performed better than HC. FG even outperformed the explorer (EXP) strategy in a dynamic environment with a moderately increasing trend. This suggests that under certain levels of feedback error and system dynamics, a decision-maker could be better off making random guesses than employing a systematic strategy.

There are also a number of instances where the performance of the FG strategy equaled that of the systematic strategies. In these situations, random guesses may also be the preferred strategy, particularly after taking into account the time and effort required to execute a systematic strategy. At low levels of feedback error, the FG strategy performed just as well as the hillclimber and at higher levels of error, FG was equivalent in performance to the EXP strategy. Also, in environments with relatively low levels of decreasing trend dynamics, the performance of the FG strategy was on par with that of the hillclimber. These findings suggest that in environments where there is no difference in performance between random guesses and systematic strategies, decision-makers could be better off choosing randomly and conserving their cognitive resources.

The random best yet strategy (RBY) proved to be remarkably robust throughout a wide range of environmental conditions. Relative to the baseline environment, the performance of RBY was unchanged at all levels of error, complexity, and dynamics evaluated in these experiments. On average, the RBY strategy performed at about ten percentage points higher than the FG strategy, making it a good "all-purpose" strategy across a wide variety of environmental conditions. Although it was not the best strategy in any one environment, it was never the worst strategy. The RBY strategy requires less cognitive effort than the systematic strategies, striking a good balance between performance, memory demands, and computational requirements.

In theory, however, the RBY strategy would suffer performance degradation at extremely high levels of error that are well beyond those evaluated in these experiments. Extreme levels of feedback error would cause frequent changes in the payoff from the best-yet settings. This leads to constantly changing best-yet settings and effectively hobbles the RBY strategy to the point where it behaves just like the FG strategy.

The HC strategy is a reasonably good performer in the baseline environment, but is vulnerable to a number of environmental factors. Even at very low levels of feedback error, performance of the HC strategy declines rapidly due to its small step search logic and the inability to separate error from the true payoff. The HC strategy is also susceptible to becoming stuck on local peaks in complex environments where there are non-linear coefficients in the production function. The HC strategy also performs sub-optimally in many of the dynamic environments evaluated in Experiment Four.

EXP is the best performing strategy in most of the environments examined in these experiments. The EXP strategy combines the small step sensitivity of the hillclimber strategy with large step logic to identify the proper starting points. However, like the hillclimber, the explorer strategy also has specific weaknesses. In particular, the EXP strategy performs rather poorly in dynamic environments with an increasing trend. At high levels of feedback error, the EXP strategy also suffers from the same problems as the HC strategy. In spite of these limitations, the EXP strategy is the best overall strategy when it is utilized in the appropriate environment.

The performance of the feedback aggregation strategy variants (RBY2, HC2, and EXP2) was less than what was expected. One possible reason for this performance is the limited number of time periods in this task. Perhaps in longer tasks, the aggregation strategies would perform better than their non-aggregating counterparts.

These experiments generated some significant results that have important implications for human decision-makers. First of all, strategies requiring more computations and cognitive effort don't always lead to better performance. It was not uncommon to have the completely random strategy (FG) and/or the random-best-yet strategy (RBY) outperform the more sophisticated systematic strategies. Potentially, human decision-makers that use a "logical" strategy like HC or EXP may be lulled into thinking that their strategy is better than random. Secondly, no single strategy performed well in all environments, not unlike the findings of Payne, Bettman, and Johnson (1993). To be effective across a variety of environmental conditions, a decision-maker should possess a repertoire of problem solving strategies and be able to know when to use them.

CHAPTER THREE

Task

The task used in the laboratory experiment was identical to that used in the simulation experiments. However, there were some contextual details added to engage the subjects in a more realistic decision-making scenario. The context was the Genesis Business Strategy Game. The specifics of the game are described in the following section.

Genesis Business Strategy Simulation Software

The task is operationalized as a computer program called Genesis that runs on an IBM personal computer network. The decision-maker is presented with information on a video display and responds interactively through keyboard commands. Figure 1 shows the computer screen display. The screen is divided into three regions. Each is discussed in more detail below.

Data entry and performance feedback section. The upper third of the display contains an area for setting the intensity level for each of the five strategic variables and information about prior period performance. In each time period, players determine the intensity level, from zero to four, for each variable: selling price, production level, R&D budget, advertising expenditure, and distribution channels. After the current period settings are made, players submit their selections and obtain the payoff (profit or loss) for that period. Two types of profit information are provided to the player: discrete time period profit (or loss) and cumulative profit (or loss).

Graphic performance data section. The center third of the computer screen provides a graphical display of prior period performance for both the player and three competitors. The competitor data were generated by the highly active decision strategies developed for the computer simulations. The three competitors are the hillclimber, random-best-yet, and the random strategy (FG). The three highly active strategies perform in the identical task environment as the subject. The reason for including competitor data in the display was to show subjects that alternative strategies produce different results. Although this introduces additional complexity, it was thought that the competitor data would motivate subjects to try harder and experiment with alternative decision strategies to improve their performance. In addition, there is a digital display of cumulative performance data for each competitor.

Instructions and error message section. The lower portion of the display is used to remind subjects of the basic rules for operating the simulation game and to flash warning messages when invalid actions are taken.

The Genesis program automatically collects detailed records of each subject's actions for each period of the task and the resultant payoffs. The output is stored in a database on the network file server. In addition, the computer program tracks and records the amount of time each subject uses to complete the task.

Subjects

Participants in the experiment were undergraduate students at the University of Arizona enrolled in business management courses. Each of the 80 subjects in the two experiments was randomly assigned to a treatment group. Subjects were given course credit for their participation. They were told in advance that the amount of course credit they would receive was directly related to their performance in the task. The average age of the participants was 22. There were approximately an equal number of males and females.

Procedure

Subjects were seated at a personal computer where they were instructed on the objectives of the task and told how to operate the controls. The instructions took the form of several computer screens of information. Subjects were told that they were participating in a study of strategic decision-making behavior in emerging markets where the traditional relationships among variables were not applicable. They were also told that the strategy aspects of the task were not intended to be a completely accurate representation of actual business strategy situations. Subjects were reminded that the objective of the simulation game was to maximize profit. Following the instructions, participants were allowed to play a practice run of 30 trials of the simulation game that was not counted in their total score. Note that the term "run" is used to describe the 30 time period task. Subjects then proceeded to play two runs of the simulation game with slight changes to the production function from run to run. The changes consisted of reordering the coefficient weights in the production function as well as re-labeling the descriptor (price, advertising, etc,) on the input display. The production function took the form: $P_t = w_1X_1 + w_2X_2 + w_3X_3 + w_4X_4 + w_5X_5 + d_t + e_t$. Subjects were told to treat each run as a unique product. The maximum possible payoff for each run remained constant. The ordering of the two products was varied to minimize any ordering effects.

Half of the subjects in each treatment group worked on the two products in A-B order, while the other half of the subjects performed the task in B-A order. At the conclusion of the two simulation runs, subjects answered some demographic questions and were debriefed on their perceptions of the task.

Hypotheses

We expected that people would begin the task by manipulating one variable at a time to see if they can figure out, or at least approximate, the relationship between the variables and the payoff. This approach can be thought of as a form of sensitivity analysis, whereby the subject attempts to understand how his or her actions affect the resultant profits. Single variable exploration can be likened to the systematic strategies used by the hillclimber and explorer in the simulation experiments. The period-to-period change in step size of the focal variable may be one level, a la the hillclimber, or four levels, similar to the explorer HAS (highly active strategy).

Further, we anticipated that if payoffs are increasing over time or if performance is better than the competition, subjects will persist in their approach of changing a single variable per time period. When subjects perceive their problem solving approach to be successful relative to prior period performance or in comparison to the competition, they are expected to continue to utilize that approach in future periods.

In contrast, if payoffs are declining over time or if subjects are under-performing the competition, people are expected to change their problem solving approach and experiment with something approximating a random strategy, i.e. changing two or more variables in a time period. When subjects perceive their approach to be unsuccessful, i.e.

profits are declining or they are uncompetitive, they are expected to abandon the sensitivity analysis methodology and choose their settings in a more random fashion, relying on luck rather than skill to maximize their profits.

The following hypotheses were developed based on the relative performance of the highly active strategies (HAS) in the simulation experiments.

We expected that the performance of human subjects will improve with experience in the task (Sterman, 1988). In the laboratory experiment, subjects will have two simulation runs, each in the same task environment but with a slightly different production function. The performance in the second run is expected to be higher than the first run due to familiarity with the task, increased proficiency in interpretation of feedback, and more consistent application of decision rules.

Hypothesis 5: Subjects' performance will improve with successive runs of the simulation within the task environment.

In the baseline task environment, the relationships between the variable settings and the payoffs are unaffected by feedback error, system dynamics or complexity. The simulation experiments showed that strategies that changed only one variable per period (HC and EXP) outperformed those strategies that changed more than one variable per period (RBY and FG) in this type of environment. We expected that human subjects would find that changing one variable at a time will lead to generally increasing payoffs and will persist in their use of that approach for the entire task. This leads to the following hypothesis:

Hypothesis 6: In environments with no error, dynamics, or complexity, subjects will manipulate one variable per time period.

In dynamic environments where the underlying trend is increasing over time, the majority of the actions taken by the decision-maker will produce payoffs that are greater than the payoff in the previous period. Since the decision-maker only sees the total payoff, he or she will be unable to distinguish between the payoff attributable to the action taken and the payoff generated by the underlying trend. We expected that decision-makers will view their approach as being successful and will persist in their strategy of manipulating a single variable per time period for the duration of the task. The simulation experiments found that, in this type of environment, strategies that change one variable per period outperform strategies that change multiple variables in a time period.

Hypothesis 7: In environments with an increasing dynamic trend, subjects will manipulate one variable per time period.

Environments with feedback error will make it difficult for subjects to identify relationships between their actions and the resultant payoffs. In fact, a subject who applies the identical settings in two different time periods is likely to observe two widely different payoffs. This may discourage the use of a single variable approach to profit maximization. Additionally, the performance of one or more competitors is likely to be greater than that of the single-variable small-step approach. It is expected that decision-makers will abandon the single variable approach in favor of a multiple variable approach in environments containing feedback error. Thus, the following hypothesis:

Hypothesis 8: In environments with feedback error, subjects will manipulate two or more variables per time period, on average, over the duration of the task.

In dynamic environments with a decreasing underlying trend, subjects will begin the task by changing one variable per time period. However, the decreasing trend will cause the majority of the subjects' actions to produce payoffs that are lower than the payoff in the prior period. In addition, subjects are likely to observe that one or more competitors are outperforming them. Under these conditions, subjects are expected to abandon their single variable approach and adopt a more random approach of changing the settings of multiple variables in a time period. The following hypothesis is proposed:

Hypothesis 9: In environments with a decreasing dynamic trend, subjects will manipulate two or more variables per time period.

Complex environments favor a methodical approach to locating the optimum settings that generate maximum profits. The non-linear nature of the payoff function rewards those approaches that systematically change a single variable per time period. It is expected that subjects will persist in their use of a single variable approach because it will produce profits that equal or exceed their competitors' performance. This is consistent with the simulation experiments, which found that systematic HAS performed better than random strategies.

Hypothesis 10: In environments with complexity, human subjects will manipulate one variable per time period.

Results of Experiment Five

In the first laboratory experiment (referred to as Experiment Five), a two by two factorial design was employed using feedback error and a positively sloped dynamic trend as treatment factors. In the error condition, the payoff function contained an error term with a standard deviation of six, normally distributed with a mean of zero. In the positive trend condition, the dynamic component of the payoff function had a slope of two.

The primary dependent variable in the task was performance, expressed as a percentage of the maximum possible payoff. Performance was analyzed in a number of different ways, including 1) between-group comparisons of both total performance and performance per time block, 2) within-group analysis of performance change between block one and block three, and 3) within-subjects comparison of performance in the first run to performance in the second run. For clarity, the results of the experiment will be analyzed in that order.

Total performance across all time periods and simulation runs was highest for subjects in the baseline condition (no error, no positive trend) at 54% of maximum (see Table 7). The positive trend treatment condition performed second best at 48%. Subjects in the error condition finished third at 45% and the treatment group with both error and positive trend performed lowest at 40%. Main effects were significant for both error ($F(79) = 8.08$, $p < .01$) and positive trend ($F(79) = 4.36$, $p < .05$), however there were no interactions between the treatment factors. Table 8 contains a summary of the univariate F statistics pooled across all simulation runs and time periods.

Performance was also assessed by dividing the task into three blocks of ten time periods each, referred to as blocks one, two, and three. Block one includes time periods one through ten, block two includes periods 11 through 20, and block three covers periods 21 through 30. In block one, performance was highest in the baseline and error conditions (45%) and lowest for subjects in the positive trend and error plus positive trend treatment conditions, at 34% of maximum. A main effect for positive trend was significant ($F(79) = 6.14$, $p < .05$). See Table 9 for a summary of univariate F statistics by time block. Performance in block two was highest in the baseline condition (56% of maximum) and lowest for subjects in the error and error plus positive trend conditions (40%). Significant main effects were present for the error treatment ($F(79) = 9.29$, $p < .01$) only. In block three, once again performance was highest in the baseline condition (61% of maximum) and lowest in the error plus positive trend condition (46%). Significant main effects were present for error ($F(79) = 7.59$, $p < .01$) and marginally significant for positive trend ($F(79) = 3.67$, $p < .10$).

The within-groups analysis of performance also produced some interesting results. Performance was compared across the three time blocks. Paired sample T-tests were used to compare performance in block one to performance in block three. The T-test statistics are summarized in Table 10. In the baseline condition, performance increased significantly between blocks one and three, moving from 45% to 61% ($t = 4.52$, $p < .01$). In the error condition, performance remained essentially unchanged between block one and block three. In the trend condition, performance increased significantly from 41% in block one to 52% in block three ($t = 2.97$, $p < .01$). In the error plus trend treatment condition, performance also increased significantly over time, rising from 34% to 46% ($t = 3.06$, $p < .01$).

A second within-groups analysis was performed to compare the group mean performance to the performance of individuals within the group. We call this approach the trajectory analysis (see Table 11), which is designed to reveal any disparity between the group mean and individual data, such as what might be caused by a bimodal distribution. The procedure for determining a subject's trajectory was as follows. The change in performance between time blocks was classified into one of two categories, up or down. If performance increased from block to block, it was classified as "up". If performance declined from block to block, it was categorized as "down". Each individual was assigned two trajectory labels, one for the change in performance between block one and block two and one for the change between blocks two and three. The analysis suggests that the group mean performance numbers are reasonably close to the performance of most of the individuals in the group. For example, in the baseline condition, the group mean trajectory is categorized as up-up since the block one to block two performance increased by 11% (up) and performance improved between blocks two and three by 5% (up). The data in the table shows that 13 of the 20 subjects in the baseline group have trajectories that match the group mean trajectory. In sum, the analysis indicates that the group mean data reflect the performance of most of the individuals within the group.

A within-subjects analysis was conducted to evaluate the change in performance between a subject's first and second runs. We hypothesized that there would be a significant performance improvement between the first and second runs. The results are summarized in Table 12. Across all treatment conditions, performance increased from 45% in the first run to 48% in the second run, however this increase was not significant ($t = 1.51$, n.s.). Thus, Hypothesis 5 was not supported. The largest increase was in the error condition where mean performance improved by 7%, moving from 42% to 49%. Although the increase was in the predicted direction, it was not statistically significant. A closer look at the data revealed that only five of the ten subjects in the error group actually increased their performance from run to run. The remaining five subjects either stayed the same or experienced a performance decline.

It is possible to gain some insight into an individual's problem-solving strategy by analyzing the process indicator variables. Process indicators include the number of variables changed from period to period, the change in intensity level of those variables, the number of periods with duplicate settings, and the number of consecutive periods with duplicate settings. These process indicators were chosen because the simulation experiments revealed that these measures were capable of differentiating the various decision strategies. It is not intended to be an exhaustive list of all of the potential process indicators. Each of the process indicators will be reviewed in detail in the following paragraphs. Analysis of the process indicator variables will be performed both between groups and within-groups.

The average number of variables per time period is calculated by comparing the intensity level of each of the five variables in one period to the intensity levels in the following period. If the intensity level is changed, regardless of the magnitude of the change, it is counted as one. If the intensity level is unchanged, it is counted as zero. For each time period, the number of changed variables can range between zero and five. The mean number of changed variables was computed for the three blocks of ten time periods as well as for the entire task. The mean values for all of the process indicator variables are reported in Table 7. The mean number of variables changed per period throughout the entire task was 1.9 and there were no significant differences among the four treatment groups. Closer inspection of individual performance reveals that of the 80 simulation runs in this experiment across all treatment conditions, only seven runs manipulated one variable or less throughout the task. It was hypothesized that subjects in the baseline condition and the upward trend condition would manipulate only one variable per time period in order to attain maximum performance. The findings indicate that nearly two variables were changed per period, thus hypotheses six and seven are not supported. Subjects in the error condition were expected to manipulate two or more variables per period to optimize performance. Based on an actual value of 1.9 changed variables per period, hypothesis eight receives moderate support.

In block one, subjects in the positive trend and error plus trend conditions changed more variables than those in other groups. A main effect was significant for the trend factor ($F(79) = 12.9$, $p < .01$). In block two, all groups ranged between 1.7 and 2.0 variables changed per period and there were no significant between-group differences. In block three, a marginally significant main effect was evident for the trend treatment ($F(79) = 3.25$, $p < .10$).

A within-groups analysis was performed on the number of variables changed from block to block using paired sample T-tests. In the baseline and error conditions there were no significant differences in the number of variables changed between blocks one and three. In the trend condition, subjects began the task by manipulating an average of 2.4 variables per period in block one, declining to 1.5 variables per period in block three. This decrease was statistically significant ($t = 3.41$, $p < .01$). In the error plus trend group the results were similar. There was a significant decrease in the number of variables changed between blocks one and three. In block one, there was an average of 2.7 variables changed per time period and in block three, the average number of variables changed per period decreased to 1.3 ($t = 8.28$, $p < .01$).

The second process indicator variable to be analyzed is the average period-to-period change in intensity level. This variable is calculated by comparing the intensity level for each variable to the intensity level for each variable in the following time period. The absolute values of the differences in intensity levels are summed and then divided by the number of variables changed to determine the change in intensity level for that time period. For example, if all five variables were set to zero in period one and all five variables were set to intensity level four in period two, it would result in a value of four. The mean change in intensity level was computed for the three blocks of ten time periods as well as for the entire task. Over the entire task, all groups averaged a change in search intensity level of between 1.3 and 1.4. There were no significant between-group differences.

In block one, subjects in the positive trend and error plus positive trend conditions made larger intensity level adjustments than other subjects did. A main effect was present for positive trend ($F(79) = 5.66$, $p < .01$). Subjects in the baseline and error conditions changed the intensity levels by an average of 1.3 and 1.2, respectively. The positive trend condition and error plus positive trend condition had mean intensity level changes of 1.5 and 1.4 per variable per period. In blocks two and three, all groups ranged between 1.2 and 1.5 levels per period and there were no significant between-group differences.

A within-groups analysis was also performed on the change in intensity level from block to block using paired sample T-tests. In the baseline and positive trend conditions, there were no significant differences in the average change in intensity level from block one to block three. In the error condition, the average change in intensity level increased significantly from 1.2 to 1.3 between blocks one and three ($t = 2.36$, $p < .05$). In the error plus positive trend condition, there was a significant decrease in search intensity, moving from 1.4 to 1.2 between blocks one and three ($t = 3.73$, $p < .01$).

The third process indicator is the number of periods with duplicate settings. With this measure, the duplicate settings do not have to be in consecutive periods to be counted. There are a maximum of nine possible duplicates in a ten period block even if the same settings were used for all ten periods. Similarly, there are a maximum of 29 possible duplicates in one run. The total number of duplicates for one run may be greater than or equal to the number of duplicates in each of the three blocks, for example when the same settings are applied in periods 5 and 15. Over all 30 periods in the task, the baseline and error treatment conditions had 11.7 and 11.4 periods with duplicate settings, respectively. Subjects in the positive trend condition had 9.8 periods with duplicate settings, while the error plus positive trend condition had 7.6 duplicates. A main effect was present for the positive trend condition ($F(79) = 4.81$, $p < .05$).

In block one, subjects in the baseline condition had a mean of 2.3 periods with duplicate settings. In the error condition, there were an average of 2.9 duplicates. Subjects in the positive trend condition averaged 1.3 duplicates in the first block, while those in the error plus positive trend condition had only 0.4 duplicates. There was a main effect for positive trend ($F(79) = 18.2$, $p < .01$) and a moderately significant interaction ($F(79) = 3.34$, $p < .10$). In the second and third blocks there were no significant between-group differences.

A within-groups analysis of the use of duplicate settings was performed using paired sample T-tests. In the baseline condition, the number of periods with duplicate settings increased significantly from 2.3 periods with duplicate settings to 4.0 duplicates between blocks one and three ($t = 2.49$, $p < .05$). Subjects in the error condition had essentially no change in the number of duplicates between blocks one and three. In the positive trend condition, the number of duplicates increased significantly from 1.3 to 3.6 between blocks one and three ($t = 3.66$, $p < .01$). The error plus positive trend condition had the largest increase, beginning the task with 0.4 duplicates in block one and increasing to 3.5 duplicates in block three ($t = 6.30$, $p < .01$).

The fourth process indicator variable is the number of consecutive periods with duplicate settings. This differs from the previous measure in that the duplicate settings must be in consecutive periods to be counted. There are a maximum of nine possible consecutive duplicates in a ten period block even if the same settings were used for all ten periods. There are a maximum of 29 possible consecutive duplicates in one run. The total number of periods with consecutive duplicates for one run may be up to two more than the number of consecutive duplicates in each of the three blocks. For example, this could occur when same settings are applied in periods 10 and 11 or in periods 20 and 21. Across the total task, subjects in the baseline condition used the highest number of consecutive duplicates with an average of 6.3 periods. The error condition had 5.6 consecutive duplicates on average, while the positive trend condition used a mean of 6.1 periods with consecutive duplicates. The error plus positive trend group had the fewest number of consecutive duplicates with 3.6. There were no significant main effects.

In block one, the baseline and error conditions had 1.8 and 1.9 consecutive duplicates, respectively. Subjects in the positive trend condition used approximately half as many consecutive duplicates (0.9) in the first block. Subjects in the error plus positive trend condition used the fewest number of consecutive duplicates (0.2). There was a main effect for positive trend ($F(79) = 13.4$, $p < .01$). In block two, the range between conditions was smaller. The baseline condition had 1.6 periods with consecutive duplicates, the error condition had 1.3, and the positive trend group had 2.2 consecutive duplicates. Subjects in the error plus positive trend condition had the fewest number of consecutive duplicates with an average of 1.0. There was a moderately significant main effect for error in block two ($F(79) = 3.18$, $P < .10$). In block three, all groups ranged between 2.1 and 2.9 consecutive duplicates and there were no significant between-group differences.

A within groups analysis was performed to assess the differences in the number of consecutive duplicates used in the first and third time blocks. Between blocks one and three, subjects in the baseline and error conditions increased the number of periods with consecutive duplicates, however the differences were not significant. In the positive trend condition, the number of consecutive duplicates increased significantly from 0.9 to 2.7 between blocks one and three ($t = 3.37$, $p < .01$). In the error plus positive trend condition, there was also a significant increase between blocks one and three ($t = 3.72$, $p < .01$). The mean number of consecutive duplicates increased from 0.2 in block one to 2.2 in block three.

A correlation matrix was created and is reported in Table 13. A number of values in the matrix indicate strong simple correlations among the variables. One of the process indicator variables, the average number of variables changed per period, has a strong negative correlation with performance ($r = -.469$, $p < .01$). This inverse relationship implies that performance is greater when the number of variables changed per period is lower. Another strong correlation exists between the number of periods with duplicate settings and the number of periods with consecutive duplicate settings ($r = .824$, $p < .01$).

A regression analysis was performed using performance as the dependent variable and the four process indicators as the independent variables. The strong simple correlations evident in the correlation matrix indicate the potential for multi-collinearity. Prior to running the regression, the process indicator variables were checked for multi-collinearity using a statistic known as the Variance Inflation Factor (Neter, Kutner, Nachtsheim, & Wasserman, 1996). The presence of serious multi-collinearity can artificially inflate the R^2 values and lead to incorrect statistical inferences about the relationships between the independent variables and the dependent variable. A Variance Inflation Factor (VIF) value of ten or more is indicative of serious multi-collinearity problems. Table 14 reports the results of the VIF analysis. The highest VIF value is 3.46, suggesting that multi-collinearity should not adversely affect the R^2 values. The results of the linear regression analysis are reported in Table 15. Standardized coefficients were used to account for the unequal variances among the variables in the regression model. When the model was created using data from all groups, the adjusted R^2 value = .32 indicating that a moderate amount of the variance in performance can be explained by the process indicator variables. The strongest predictors of performance include 1) the number of variables changed per period, 2) the change in intensity level per variable per period, and 3) the number of periods with consecutive duplicate settings. The results of the regression analysis indicates the first two variables are significant at $p < .01$ and the number of consecutive duplicates is significant at $p < .05$.

Results of Experiment Six

The second laboratory experiment (referred to as Experiment Six) was a two by two design that evaluated the effects of a negatively sloped dynamic trend and complexity on task performance. In the negative trend condition, the dynamic component of the payoff function had a slope of minus two. Complexity was operationalized as a non-linear coefficient applied to one of the five action variables.

As in the previous experiment, the primary dependent variable in the task was performance, expressed as a percentage of the maximum possible payoff. All aspects of the task except the task factors of negative trend and complexity were identical to those in Experiment Five. The results are also summarized using the same approach. First is the between groups analysis, followed by the within groups analysis, and concluding with the within subjects analysis.

Total performance across all time periods and simulation runs was highest for subjects in the baseline condition (no trend, no complexity) at 64% of maximum (see Table 16). Subjects in the complexity condition performed second best at 57%. Subjects in the negative trend condition finished at 53% and the group with both trend and complexity performed at 52%. Main effects were significant for the negative trend ($F(89) = 11.77$, $p < .01$) but not significant for complexity ($F(89) = 2.50$, $p < .11$). There were no interactions between the treatment factors. Table 17 contains a summary of the univariate F statistics pooled across all simulation runs and time periods.

Performance was also assessed by time block. In block one, all groups performed between 48% and 53% of maximum and there were no significant between-group differences. See Table 17 for a summary of univariate F statistics by time block. Performance in block two was highest for subjects in the baseline condition at 67%. This was followed by a 58% performance for subjects in the complexity condition. Participants in the negative trend and complexity plus trend conditions performed at 54% and 50%, respectively. Significant main effects were present for both the negative trend condition ($F(89) = 16.1$, $p < .01$) and complexity ($F(89) = 6.2$, $p < .05$). In block three, subjects in the baseline condition again performed best at 71%. The subjects in the complexity condition performed at 64%. The negative trend and negative trend plus complexity groups tied for the lowest performance in block three at 54% of maximum. There was a significant main effect for negative trend ($F(89) = 21.1$, $p < .01$).

The within-groups analysis of performance also produced interesting results. As in the prior analysis, performance was compared using three blocks of ten time periods each. Paired sample T-tests were used to compare performance in block one to performance in block three. The T-test statistics are summarized in Table 19. In the baseline condition, performance increased significantly between blocks one and three, moving from 53% to 71% ($t = 4.71$, $p < .01$). In the negative trend condition, there was a marginally significant difference between block one and block three ($t = 1.74$, $p < .10$), increasing from 50% to 54%. In the complexity condition, performance in block one was 48% and increased to 64% in block three ($t = 5.80$, $p < .01$). In the negative trend plus complexity condition, performance increased slightly, from 51% to 54% (n.s.) between blocks one and three.

The trajectory analysis, using the same procedure as described in the results section of Experiment Five, revealed no important differences between group mean performance and the performance of the majority of the individuals within each treatment group. The results of the trajectory analysis are reported in Table 20.

A within-subjects analysis was conducted to evaluate the change in performance between a subject's first simulation run and his or her second simulation run. It was hypothesized that there would be a significant performance improvement between the first and second runs. The results are summarized in Table 21. Across all treatment groups, performance decreased from 57% in the first run to 56% in the second run, however the decrease was not significant. This was consistent with the findings in Experiment Five, and Hypothesis Five was not supported. The only performance increase was in the baseline condition where mean performance improved by 4%, moving from 62% to 66%. Although the increase was in the predicted direction, it was not statistically significant.

An analysis was conducted on subjects' behavior using the process indicator variables. Process indicators include the number of variables changed from period to period, the change in intensity level of those variables, the number of periods with duplicate settings, and the number of consecutive periods with duplicate settings. Each of the process indicators will be reviewed in detail in the following paragraphs. Descriptions of the procedures used to calculate the process indicator variables were included in the results section of Experiment Five and are not restated here. The analysis was performed both between groups and within-groups.

The mean number of changed variables was computed for the three blocks of ten time periods as well as for the entire task. The results can be found in Table 16. The mean number of variables changed per period throughout the entire task was 1.9. A significant difference was evident for negative trend ($F(89) = 25.5$, $p < .01$), while complexity was marginally significant ($F(89) = 2.68$, $p < .10$). Subjects in the negative trend condition were expected to manipulate two or more variables per period to optimize performance. Based on an actual value of 2.2 changed variables per period, Hypothesis

Nine is supported. It was hypothesized that subjects in the baseline condition and the complexity condition would manipulate only one variable per time period in order to attain maximum performance. The findings indicate that between 1.3 and 1.7 variables were changed per period, thus Hypothesis Ten is not supported.

In block one, there were no significant between group differences in the average number of variables changed per period. Overall, subjects changed between 1.8 and 2.2 variables per period. In blocks two and three, there was a significant difference for negative trend, ($F(89) = 29.0$, $p < .01$) and ($F(89) = 40.2$, $p < .01$), respectively.

A within-groups analysis was also performed on the number of variables changed from block to block using paired sample T-tests. In the baseline condition, the number of changed variables declined significantly from 1.8 in block one to 1.1 in block three ($t = 3.30$, $p < .01$). In the negative trend condition there was a significant increase in the number of variables changed between blocks one and three, starting at 1.9 and moving to 2.4 ($t = 2.27$, $p < .05$). In the complexity environment, subjects changed a mean of 2.0 variables in the first block, declining significantly to 1.4 variables in block three ($t = 4.88$, $p < .01$). In the negative trend plus complexity condition, subjects began the task by manipulating 2.2 variables in block one, increasing marginally to 2.4 variables per period in block three ($t = 1.99$, $p < .10$).

The second process indicator variable to be analyzed is the average period-to-period change in intensity level. The mean change in intensity level was computed for the three blocks of ten time periods as well as for the entire task. Over the entire task, subjects in the baseline environment adjusted each changed variable by an average of 1.6

intensity levels, while those in the complexity condition averaged 1.1 levels ($F(89) = 12.1$, $p < .01$). There were no other significant between-group differences.

In blocks one, two, and three the only significant between-group differences were in the complexity condition. Subjects had significantly lower average changes in the intensity levels compared to the other groups.

A within-groups analysis was also performed on the change in intensity level from block to block using paired sample T-tests. The only group to change significantly over time was the negative trend group which increased from 1.3 in block one to 1.6 in block three ($t = 2.68$, $p < .05$). There were no other within-group changes across the three time blocks.

The third process indicator is the number of periods with duplicate settings. Over all 30 periods in the task, subjects in the baseline and complexity conditions had 16.4 and 13.8 periods with duplicate settings, respectively. Subjects in the negative trend condition had 11.0 periods with duplicate settings, while the negative trend plus complexity group had 8.5 duplicates. The ANOVA revealed a main effect for both negative trend ($F(89) = 26.2$, $p < .01$) and complexity ($F(89) = 5.8$, $p < .05$) with no interaction.

In block one, subjects in the baseline condition had a mean of 3.4 periods with duplicate settings. In the negative trend and complexity conditions, subjects averaged 2.2 and 2.3 duplicates, respectively. Subjects in the negative trend plus complexity group averaged 1.6 duplicates in the first block. There was a main effect for trend ($F(89) = 5.8$, $p < .05$) and for complexity ($F(89) = 4.5$, $p < .05$). In the second block, subjects in the negative trend condition used significantly fewer duplicates ($F(89) = 41.8$, $p < .01$). In block three, subjects in the negative trend and complexity conditions had fewer periods with duplicate settings. Main effects were evident for both negative trend ($F(89) = 54.4$, $p < .01$) and complexity ($F(89) = 4.8$, $p < .05$).

The within-groups analysis was performed using paired sample T-tests. In the baseline condition, the number of periods with duplicate settings increased significantly from 3.4 duplicates to 5.7 duplicates between blocks one and three ($t = 3.04$, $p < .01$). In the complexity group, the number of duplicates increased significantly from 2.3 to 4.5 between blocks one and three ($t = 3.85$, $p < .01$). Subjects in the negative trend and negative trend plus complexity conditions had essentially no change in the number of duplicates between blocks one and three.

The fourth process indicator variable is the number of consecutive periods with duplicate settings. Across the entire task, subjects in the baseline condition used the most, with an average of 10.8 periods with consecutive duplicates. Subjects in the complexity condition had 7.6 consecutive duplicates on average, while the negative trend and negative trend plus complexity groups used an average of 2.0 and 1.5 consecutive duplicates, respectively. A significant main effect was found for negative trend ($F(89) = 61.3$, $p < .01$) and complexity was marginally significant ($F(89) = 3.14$, $p < .10$).

In block one, the baseline and complexity groups had 2.4 and 1.2 consecutive duplicates, respectively. The negative trend group used 0.9 consecutive duplicates on average, while subjects in the trend plus complexity group used 0.5. The ANOVA revealed a main effect for negative trend ($F(89) = 8.6$, $p < .01$) and complexity ($F(89) =$

4.4, $p < .05$). In blocks two and three, significant main effects were found for negative trend only, ($F(89) = 34.1$, $p < .01$) and ($F(89) = 68.6$, $p < .01$), respectively.

A within groups analysis was performed to assess the differences between the first and third time blocks. Between blocks one and three, subjects in the baseline condition significantly increased the number of periods with consecutive duplicates from 2.4 to 4.4 ($t = 2.27$, $p < .05$). Subjects in the negative trend condition decreased the number of consecutive duplicates used from 0.9 to 0.4, which was marginally significant ($t = 1.98$, $p < .10$). Subjects in the complexity condition had a significant increase in the use of consecutive duplicates, moving from 1.2 in block one to 3.4 in block three ($t = 3.67$, $p < .01$). The negative trend plus complexity group was essentially unchanged across time blocks.

A correlation matrix was created and is reported in Table 22. A number of values in the matrix indicate a strong correlation among the variables. All variables but the average period-to-period change in intensity level have a moderately strong bivariate correlation with performance. The strongest relationship with performance is the number of variables changed per period, which is inversely related ($r = -.571$, $p < .01$), consistent with the findings in the first experiment. Also consistent with the finding in the first experiment is the strong correlation between the number of duplicates and the number of consecutive duplicates ($r = .810$, $p < .01$).

A regression analysis was performed using performance as the dependent variable and the four process indicators as the independent variables. As in Experiment Five, the process indicator variables were checked for multi-collinearity using the Variance

Inflation Factor statistic (Neter, Kutner, Nachtsheim, & Wasserman, 1996). Table 23 reports the results of the VIF analysis. The highest value in the table is 3.70, suggesting that multi-collinearity should not adversely affect the R^2 values. The results of the linear regression analysis are reported in Table 24. Standardized coefficients were used to account for the unequal variances among the variables in the regression model. When the model was created using data from all groups, the adjusted R^2 value was .34, indicating that a moderate amount of the variance in performance can be explained by the process indicator variables. Across all groups, the strongest predictor of performance is the average number of variables changed per period ($p < .01$).

Discussion

This section will begin with a summary of the key findings of the two laboratory experiments and conclude with a comparison of the performance of human subjects to the random benchmark strategy from the simulation experiments. In Chapter Four, a theory of human problem solving behavior in novel task environments will be presented.

In the first laboratory experiment (Experiment Five), the effects of feedback error and positive trend were evaluated. Each of the treatment factors resulted in lower performance relative to the baseline (no error or positive trend) environment. Performance was lowest for subjects in the error plus positive trend condition. Over time, performance increased for all subjects except for those in the error condition. This finding suggests that subjects in the baseline and positive trend environments are able to learn about the task and improve their performance, while those in the error environment are not. Inaccuracies in performance feedback limit the ability of subjects to learn about their environment and improve their performance.

There were also differences in the process indicators. Subjects in the positive trend condition had a much greater search radius (average number of variables changed per period) early in the task when the effects of the trend manipulation generated negative profits. However, over time as profits became positive, the search radius of subjects in the positive trend condition was equal to that of subjects in the other groups. Search stability (use of duplicate settings) was highest for subjects in the baseline condition. As the task progressed, subjects in the baseline and positive trend environments decreased their search radius and increased their search stability, while subjects in the error condition had no change in their search radius or search stability over time.

Surprisingly, there was no within-subjects performance improvement from run to run. It was expected that performance would improve between an individual's first run and their second run. Although the two tasks were very similar to one another, the evidence suggests that the differences did not allow the subjects to utilize any of the knowledge gained in the first run to improve their performance in the second run.

In the second laboratory experiment (Experiment Six), the effects of negative trend and complexity were assessed. Again, performance was highest in the baseline condition, followed closely by the complexity group. It is expected that a higher level of complexity than what was actually used would have resulted in significantly lower performance. Subjects in the negative trend condition had the worst total performance. Over time, performance improved for subjects in the baseline and complexity conditions, but not for subjects in the negative trend condition. This suggests that learning takes place in the baseline and complexity conditions, but is blocked when the environment is characterized by negative trend.

Analysis of the process indicators revealed some interesting findings. Subjects in the negative trend condition had a higher search radius and lower search stability, particularly in the latter periods of the task when profits were negative and the underlying trend was declining. Subjects in the complexity condition had a lower level of search intensity (average change in intensity level settings per variable changed). The reduction in search intensity occurred early in the task and stayed lower for the remainder of the game.

Again, there was no performance improvement between a subject's first and second run, consistent with Experiment Five but contrary to the hypothesis. The results suggest that there was no learning effect that enabled subjects to transfer knowledge from the first run to the second run. This finding differs from the results of experiments by Sterman (1987) and Brehmer (1992) where repeated play produced increasingly higher performance.

Another interesting finding that was observed in both of the laboratory experiments was the relatively high number of runs that performed at a level that was equal to or less than the random strategy. The random strategy chooses the settings for each task variable entirely at random without regard to prior performance and serves as a lower benchmark for performance comparisons. In total, 20 of the 170 runs (12%) were below the performance level that could have been achieved simply by choosing settings at random. The results are reported in Table 25. There were an additional six runs (4%) that were equivalent in performance to the random strategy. It is also worth noting that only two individuals had both of their runs below or equal to the random strategy performance level. The balance of the runs were generated by 22 separate individuals. The results suggest that subjects performing at this level could have saved a considerable amount of cognitive effort by simply choosing their actions at random.

In summary, the two laboratory experiments produced some interesting as well as surprising results. Many of the hypotheses about human behavior that were drawn from the simulation experiments were unsupported. For example, in the early periods of the task, subjects did not manipulate one variable per time period in the same way that simulated subjects behaved. Rather, it is theorized that they utilized a two-stage approach. The early periods in the task were used to explore the environment in order to obtain a broad, yet minimal understanding of the task before advancing to the next stage. In stage two, human subjects behaved more like their simulated counterparts in terms of low search radius and search intensity. This will be discussed in more detail as part of the problem-solving model to be discussed in the following chapter. Another surprise result was the failure of human subjects to improve their performance across similar tasks. Additional experiments may need to be run to more fully understand how and why this occurred. Finally, a relatively high number of subjects expended a fair amount of time and energy pursuing strategies that performed at a level equal to or below that of taking actions entirely at random. This finding may open the door to future work to understand why people adopt flawed problem solving strategies and persist in using those strategies for the entire task even with the knowledge that better strategies do exist.

CHAPTER FOUR

Feedback-Based Theories of Learning

This chapter will begin with a brief overview of the leading theories of learning and performance in dynamic task environments. Next is a short analysis of the barriers to learning and performance. The chapter will conclude with the presentation of a two-stage model of behavior in novel task environments.

In unknown environments, the decision-maker begins with virtually no knowledge or understanding of the world. In order to acquire knowledge and learn about the environment, action must be taken. There are no other opportunities to test, observe, or otherwise gain knowledge about the environment other than to act. Decisions are made, feedback is received, and the new information is used to make the next decision. Feedback-based accounts of behavior and learning are now well accepted in most of the social and management sciences. Forrester (1961) asserted that all decisions, and therefore learning, take place in the context of feedback loops. Figure 2 shows a simple feedback loop model where the environment provides information in the form of feedback about the effectiveness of previous decisions. The feedback is then used to guide the decision-maker in subsequent choices. The feedback loop model in Figure 2 describes the most basic type of learning. Decision-makers compare quantitative and qualitative information about the state of the environment to their goals, perceive discrepancies between desired and actual states, and take actions that they believe will move them toward the desired goal (Forrester, 1961).

However, the feedback loop model neglects an important aspect of learning. Information feedback is not the only input into the decision process. Decisions are the result of applying a decision rule to the feedback received (Argyris, 1985). The decision rules are developed by the decision-maker, subject to influences from the organization, culture, and prior experience. Building on the work of Forrester (1961), Argyris (1985) developed the single loop learning model. This model is contained in Figure 3. The single loop learning model features a decision strategy supported by a set of decision rules. The decision strategy is driven by an individual's mental model of the real world environment. The mental model contains a person's beliefs about the cause and effect relationships of their actions. In the single loop learning model, the mental model and decision rules remain static. That is, the mental model and therefore the decision strategy, do not change as a result of new information about the environment. The decision rules that are in place at the beginning of a task are left unchanged for the duration. As a consequence, single loop learning does not alter an individual's understanding of the causal structure of the environment.

The single loop model does not allow the decision-maker to update his or her mental model and decision strategy in light of new or unexpected performance feedback from the environment. Argyris' (1985) double loop learning model addresses this deficiency. An illustration of the double loop learning model can be found in Figure 4. Here, feedback about performance in the real world is used as both input into future decisions and to update an individual's mental model of the environment. The perceived cause and effect relationships between actions and outcomes are actively changing as new feedback is received. When mental models change, new decision rules are created and old rules may be modified or discarded. Thus, over time, the same information feedback may generate different decisions because the information is being processed using new decision rules.

Barriers To Learning

There are numerous obstacles to learning and performance in even the simplest of tasks. For learning to occur, the feedback loop must operate effectively. Decision-makers must be able to cycle around the loop at a faster rate than the rate of change in the environment. As the rate of change in the environment increases, the decision-maker's knowledge of the environment becomes obsolete more quickly (Sterman, 1994). Barriers to learning can be loosely grouped into two categories, the structural attributes of the environment and the limitations of human cognition. The structural attributes of the environment include feedback error, complexity, feedback delays, and dynamism. Limitations of human cognition include the choice of flawed decision strategies, the inability to properly execute decision rules, and poor inferential skills.

When feedback is incorrect or ambiguous, it can lead to misunderstanding and misinterpretation of action-outcome linkages. Feedback error distorts the quality of the performance information and is indistinguishable from the true payoff resulting from the actions taken by the decision-maker. As a result, decision-makers are unsure of the direct consequences of their actions. Even a small amount of feedback error in the environment makes learning difficult. Moderate to high levels of error make learning close to impossible. Complexity in the environment slows the learning process when the number of variables that might affect performance overwhelms the data available to rule out alternative interpretations (Sterman, 1994). Interactions among variables, non-linear relationships, and the sheer number of variables all contribute to the complexity of the task. In general, learning in a complex environment requires more thought, more experimentation, and more time than learning in environments that are comparatively less complex. Delays in receiving feedback also slow the speed of the feedback loop. Time delays between taking an action and knowing its effects are common in the real world. For example, the performance of financial investments and newly hired employees are not known for months or even years after those decisions are made. Delayed feedback slows the ability to accumulate experience, test hypotheses, learn about the environment, and improve performance. Dynamic environments hinder the learning process by confounding the direct consequences of the actions taken by the decision-maker with the underlying trend. The decision-maker is forced to make inferences about the effects of the dynamics with limited data. The degree of turbulence present in the environment will affect the ability of the decision-maker to learn and make correct (or even approximate) inferences about the underlying trend.

Limitations of human cognition have been widely studied by psychologists and others for many years. No attempt is being made for coverage of the literature on judgmental errors, biases, and bounded rationality. For an overview, see Hogarth (1981, 1987) and Davis and Hogarth (1992). Rather, this subsection will focus on a few of the potential shortcomings of human decision-makers in the types of tasks studied in this research. Initially, the decision-maker may select a decision strategy that is inappropriate for the task. There are several reasons why an improper strategy could be chosen. Perhaps the decision-maker selects a strategy that has worked well in the past in other types of environments. It is also possible that the chosen strategy is the only strategy known to the strategist. Additionally, structural attributes of the environment such as feedback error and dynamism may lead to the choice of an ineffective strategy. When any of these conditions exist, the decision-maker will begin the task with a strategy that does not fit well with the environment. Another barrier to good performance is the failure to properly execute the chosen decision strategy. Even if a good decision strategy for the environment was selected, the decision-maker may fail to properly execute the decision rules that underlie the strategy. Mental fatigue, carelessness, and inattention all contribute to strategy implementation failure. The third obstacle to learning and performance to be discussed here is the potential for the decision-maker to make incorrect inferences about the environment. When this occurs, decision-makers may fail to recognize when it is time to abandon a failing decision strategy and adopt a different approach. Alternatively, it is possible that incorrect inferences may lead to discarding an effective decision strategy. Inferential skills include the ability to detect correlations among variables, update beliefs according to Bayes' rule, and apply basic statistical concepts such as regression to the mean.

In summary, the ability of decision-makers to learn and perform in dynamic task environments is a function of the features of the environment and the ability of decision-makers to select and execute appropriate decision strategies. Successful decision-makers are those who are able to correctly match their strategy to the environment.

The Need for a New Theory

Existing models of learning and performance in dynamic environments are helpful in understanding human behavior over a broad range of generic tasks. However, in an effort to generalize the theory to a variety of tasks and contexts, the existing models are both overly complicated and lacking specificity. After reviewing the existing theories, the concept of a mental model is still unclear. What exactly is a mental model, how is it structured, and how does it change? The answers are vague and difficult to understand. Sterman (1994) defines mental models as the "sensory and cognitive structure" of the real world as it is perceived. Argyris (1985) describes mental models as an individual's beliefs about the cause and effect relationships of their actions. Sterman (1994) himself admits that the concept of mental models is somewhat vague when he says, "Most people do not appreciate the ubiquity and invisibility of mental models…". Perhaps in some tasks, mental models are useful in capturing the mechanisms used to add, modify, and delete decision rules, however in the relatively simple tasks studied in this research, the concept of mental models seems superfluous. Another limitation of the current theories is the lack of explicitly stated decision rules. The findings from the experiments reported in this paper enable a high degree of specificity of the decision rules, something absent in the generic models of feedback-based learning. The decision rules will be explained in more detail later in this chapter. Another reason for developing a new model of learning is the need to include a provision for the two-stage problem-solving behavior observed in the laboratory experiments. The phase shift from Stage One to Stage Two is not readily explained by the current theories. The specifics of the two-stage model will be covered in the following section. In summary, when applied to the types of tasks reported here, existing theories of learning are overly complicated and fall short of explaining all of the behavior observed in these experiments. However, they do serve as a useful departure point for a new theory of feedback-based learning to be presented in the following section.

Two-Stage Model of Behavior in Novel Task Environments

The results of the laboratory experiments reported earlier revealed an interesting pattern of behavior not previously reported in the literature on dynamic decision-making. It was expected that human subjects would behave in a way that was similar to one of the seven highly active strategies created for the simulation experiments. The results were surprising. Many subjects employed a two-stage approach to the problem. The initial periods in the task (Stage One) were used to learn about the environment and gain a broad understanding of the task at hand. When the transition from Stage One to Stage Two did occur, it typically happened somewhere between one-quarter and halfway through the task, depending upon the features of the task environment and the individual decision-maker. The behavior in the later periods of the task (Stage Two) was similar to the highly active strategies (HAS) used by the simulated subjects.

At the center of the new theory is a model of a two-stage process for problem solving (Figures 5 and 6). Behavior in Stage One is characterized by bold actions that involve changing multiple variables in a single time period while simultaneously making relatively large increases or decreases in the intensity level settings. In other words, people used a large search radius and applied a high level of search intensity. This type of behavior can be thought of as assessing or sizing up the environment. The objective in Stage One is to learn, in a very broad sense, how the payoffs are dispersed throughout the task environment. In this stage, subjects are looking for the answer to the question "What type of game am I playing?". Does increasing the intensity level of every variable result in higher profits (Every Little Bit Helps)? Do changes to some variables increase the profits while changes to others result in a decrease (Chutes And Ladders)? Are there interactions among variables (Tricky Mixtures)? Or is there a single combination of settings that contains all the profits (Needle In A Haystack)? At the elemental level, subjects may ask themselves questions like "Does this lever work? Do the profits go up or down when I move it? If so, by how much?".

Figure 5 contains a model of the problem solving behavior in Stage One. The cycle begins when the decision-maker takes the first action. At the beginning of the task, the decision-maker has absolutely no knowledge of the environment. Shortly after the action is taken, performance feedback is received. The performance feedback provides information to the decision-maker that is used to update his or her beliefs about the environment. This information accumulates as the task progresses, adding to the beliefs that the decision-maker holds about the distribution of payoffs in the environment. The feedback is also used to make performance comparisons. The data collected in these experiments suggest that there are four different types of performance comparisons. First, is an assessment of the trend in performance. The slope of recent performance is classified as increasing, decreasing, or zero. Second, recent performance is compared to the break-even point. Profits are categorized as either positive or negative. Third, is an evaluation of the stability of the environment. In stable environments, the application of the same settings in two different time periods will produce identical payoffs. In unstable environments, such as those that contain error or an underlying trend, applying duplicate settings will generate different payoffs. The fourth and final comparison is linearity. Here, the decision-maker assesses whether there is a linear relationship between the intensity level settings and profits for one or more of the variables.

After the performance comparisons are made, this information is used as an input into the decision rules. There are two sets of decision rules. One set governs the search radius decision and the second set of rules applies to the search intensity decision. The search radius decision begins with the question "Is recent performance producing losses?". If the performance comparison had a negative sign, then the answer to this question is yes. According to the rule, a yes response requires an increase to the search radius. The second rule in the search radius decision asks the question "Is recent performance declining?". If the slope of recent performance is negative, then the answer to this question is yes. A yes response results in yet another increase to the search radius and a return to the action-feedback loop to take the next action. A no response (slope is increasing or zero) advances the decision-maker to the next rule. The third search radius decision rule uses the stability comparison. If the environment is stable, the search radius is decreased and the decision-maker moves to Stage Two. If the environment is unstable, the search radius is maintained at the current level and the next action is taken. The search intensity decision runs in parallel with the search radius decision. The rule for the search intensity decision is straightforward. If one or more of the variables are nonlinear, i.e. the linearity comparison yields a no, the search intensity level is decreased and the next action is taken. Otherwise, the search intensity level is maintained at the current level and the decision-maker returns to the action-feedback loop.

Stage One continues until the decision-maker advances to Stage Two per the decision rule or until time expires. The transition from Stage One to Stage Two only takes place when certain conditions are met. The advancement to Stage Two requires recent performance to be positive and increasing or positive and holding steady, as well as a stable environment. If any of these conditions are not met, there will not be a shift to Stage Two. The transition to Stage Two, when it does occur, happens rather abruptly. The transition is marked by a decrease in the search radius to approximately one, where it remains for the balance of the task. As noted earlier, the transition between the two stages typically occurs between one-quarter and halfway through the task. These temporal boundaries are important. The transition rarely occurs sooner than one-quarter of the way through the task, since that is the minimum length of time required to systematically evaluate the impact of each variable on profitability and test for stability. For example, the following sequence of settings requires eight periods, equivalent to just over 25% of the available time: 00000, 40000, 04000, 00400, 00040, 00004, 44444, 00000. The upper boundary for the transition is a function of the horizon effect. People compare the time elapsed to the time remaining and conclude that it is time to start focusing on profit maximization or "harvesting". Interestingly, it was observed that people who used a more structured exploration technique in the early periods of the task made the transition to Stage Two earlier than those who used a less structured approach.

The behavior in Stage Two is more tactical in nature, in contrast to the higher-level decision strategies seen in Stage One. In Stage Two, there are fewer performance comparisons, fewer decision rules, and the rules are simpler to apply. The behavior in Stage Two is similar to the process used by the hillclimber HAS (highly active strategy). Often, the initial settings for Stage Two are those that generated the highest profits in Stage One. In other cases, the decision-maker begins Stage Two with the ending settings from Stage One. The model of Stage Two can be found in Figure 6. Stage Two uses a simplified action-feedback loop. As in Stage One, action is taken and feedback is received. The feedback is used as an input to the performance comparison. The single performance comparison is the slope of the decision-maker's recent performance. The result of this comparison is a judgment as to whether recent performance is increasing, decreasing, or holding steady. That comparison, along with the performance feedback, comprises the inputs into the decision rule. The decision rule begins with the question "Is performance improving?". If the answer is yes, then the active variable is incremented by one intensity level and the next action is taken. If the answer is no, the active variable is reset to its previous intensity level and a new active variable is selected. The new active variable is incremented by one intensity level and the next action is taken.

An interesting pattern of behavior was observed in a sizable number of subjects about mid-way through Stage Two. These subjects would suddenly change the intensity level of all five variables in one time period and then resume the single variable, single intensity level tactics from that point forward. In effect, these people are making sure that they are not stuck on a local peak where they won't be able to achieve maximum profitability.

As noted earlier, there are some environments where the transition from Stage One to Stage Two never occurs. In these worlds, people persist in the typical Stage One pattern of behavior for the entire task. For example, environments with moderate levels of error do not encourage decision-makers to advance to Stage Two. Random feedback errors mean that when the same settings are applied in two different periods they will generate different payoffs. The decision-maker interprets this condition as environmental instability and consequently there is no shift to the typical Stage Two pattern of behavior. Another type of environment where people use a large search radius and make large changes in search intensity over the entire task is the decreasing trend environment. Here, the decision-maker is faced with negative profits (losses) during the last one-half to one-third of the task. The unexpected losses are interpreted by the decision-maker as being the result of their actions. When this occurs, people begin choosing settings apparently at random in an attempt to find a combination of settings that can restore them to profitability.

In summary, this new two-stage model of behavior in novel task environments both extends and simplifies the existing theories of learning. The concept of mental models is discarded in favor of an explicit set of simple decision rules driven by performance feedback and simple performance comparisons. Stage One can be thought of as learning about the payoff characteristics of the environment. A large search radius and large intensity level changes characterize behavior in Stage One. Stage Two is like problem solving in the more traditional sense. Both the search radius and search intensity are considerably smaller than Stage One levels as decision-makers work toward maximizing the payoffs. The shift from Stage One to Stage Two occurs after the decision-maker has surveyed the environment and attained a minimal level of understanding about the payoff structure. Interestingly, in environments with feedback error or a decreasing underlying trend, the transition to Stage Two never takes place. This simple behavioral model forms the basis for further empirical testing and additional theory development in other types of novel task environments.

CHAPTER FIVE

Summary Of The First Four Chapters

Chapter One reviewed and summarized the relevant literature. The prior research was divided into three major categories including dynamic decision-making, strategy and organization theory, and fast and frugal heuristics. The existing literature was used to provide background, stimulate ideas, and identify potential directions for the dissertation. The relevant variables for study in the dissertation were identified as feedback error, task complexity, and system dynamics. Feedback error is inherent in many dynamic decision-making (DDM) tasks but has received comparatively little attention from researchers as an independent variable (Brehmer, 1992). The presence of even small amounts of feedback error leads to variation in the observability of the system, which in turn may lead to performance degradation in the task. The importance of complexity as a variable for further study is clear in both the DDM and the strategy literature. Brehmer (1992) evaluated complexity in an advanced version of the fire-fighting task where subjects were confronted with multiple conflicting goals. In the strategy literature, complexity has been found to be one of the three major environmental factors affecting organization performance (Eisenhardt, 1992). Generally, it has been found that the greater the number of variables and/or goals in the environment, the lower the performance. System dynamics (Funke, 1993, 1995), can be thought of as the degree of volatility or instability in the system, independent of the prior actions of the player. Higher levels of system dynamics have led to significant decreases in overall task performance. At the organization level, dynamics is a key determinant of performance (Eisenhardt, 1992).

Aldrich (1979) found that both the rate of change and the amount of change affect organization performance.

Chapter Two reported the results of four simulation experiments. A computerized simulation program was developed to evaluate the effects of feedback error, positive and negative trends, and complexity on the performance of seven simulated subjects. The simulated subjects were constructed using different sets of decision rules. Performance was evaluated for each of the seven highly active problem-solving strategies in a variety of task environments. The simulation experiments produced some interesting and surprising results. One somewhat counterintuitive observation is that under certain conditions it is possible for systematic strategies to perform worse than the random strategy. This was true in environments with feedback error and in dynamic, fast-changing environments. At moderate to high levels of feedback error, the random strategy clearly outperformed the hillclimber. In dynamic environments with a moderately decreasing trend, the random strategy again performed better than the hillclimber. The random strategy even outperformed the explorer strategy in a dynamic environment with a moderately increasing trend. This suggests that under certain levels of feedback error and system dynamics, a decision-maker could be better off making random guesses than employing a systematic strategy. There are also a number of instances where the performance of the random strategy equaled that of the systematic strategies. In these situations, random guesses may also be the preferred strategy, particularly after taking into account the time and effort required to execute a systematic strategy. The random best yet strategy (RBY) proved to be remarkably robust throughout a wide range of environmental conditions. Relative to the baseline environment, the performance of RBY was unchanged at all levels of error, complexity, and dynamics evaluated in these experiments. On average, the RBY strategy performed at about ten percentage points higher than the random strategy, making it a good "all-purpose" strategy across a wide variety of environmental conditions. Although it was not the best strategy in any one environment, it was never the worst strategy. The RBY strategy requires less cognitive effort than the systematic strategies, striking a good balance between performance, memory demands, and computational requirements.

The hillclimber strategy was a reasonably good performer in the baseline environment, but is vulnerable to a number of environmental factors. Even at very low levels of feedback error, performance of the hillclimber strategy declines rapidly due to its small step search logic and the inability to separate error from the true payoff. The hillclimber strategy is also susceptible to becoming stuck on local peaks in complex environments where there are non-linear coefficients in the production function. The explorer strategy is the best performing strategy in most of the environments examined in these experiments. The explorer strategy combines the small step sensitivity of the hillclimber strategy with large step logic to identify the proper starting points. However, like the hillclimber, the explorer strategy also has specific weaknesses. In particular, the explorer strategy performed rather poorly in dynamic environments with an increasing trend. At high levels of feedback error, the explorer strategy also suffers from the same problems as the hillclimber strategy.

The simulation experiments generated some significant results that have important implications for human decision-makers. First of all, strategies requiring more computations and cognitive effort don't always lead to better performance. It was not uncommon to have the completely random strategy and the random-best-yet strategy outperform the more sophisticated systematic strategies. Secondly, no single strategy performed well in all environments. To be effective across a variety of environmental conditions, a decision-maker should possess a repertoire of problem solving strategies and be able to know when to use them.

Chapter Three assessed the performance of human subjects using the same task environments as the computer simulation. In the first laboratory experiment, the effects of feedback error and positive trend were evaluated. Each of the treatment factors resulted in lower performance relative to the baseline environment. Over time, performance increased for all subjects except for those in the error condition. This finding suggests that subjects in the baseline and positive trend environments are able to learn about the task and improve their performance, while those in the error environment are not. Subjects in the positive trend condition had a much greater search radius (average number of variables changed per period) early in the task when the effects of the trend manipulation generated negative profits. Over time, as profits became positive, the search radius of subjects in the positive trend condition was equal to that of subjects in the other groups. As the task progressed, subjects in the baseline and positive trend environments decreased their search radius and increased their search stability, while subjects in the error condition had no change in their search radius or search stability over time.

In the second laboratory experiment, the effects of negative trend and complexity were assessed. Again, performance was highest in the baseline condition, followed closely by the complexity group. Subjects in the negative trend condition had the worst total performance. Over time, performance improved for subjects in the baseline and complexity conditions but not for subjects in the negative trend condition. This suggests that learning takes place in the baseline and complexity conditions but is blocked when the environment is characterized by negative trend. Subjects in the negative trend condition had a higher search radius and lower search stability, particularly in the latter periods of the task when profits were negative and the underlying trend was declining. Subjects in the complexity condition had a lower level of search intensity (average change in intensity level settings per variable changed). The reduction in search intensity occurred early in the task and stayed lower for the remainder of the game.

In both experiments, there was no performance improvement between a subject's first and second run. The results suggest that there was no learning effect between runs. Another interesting finding that was observed in both of the laboratory experiments was the relatively high number of runs that performed at a level that was equal to or less than the random strategy. 12% of the runs were below the performance level that could have been achieved simply by choosing settings at random. The results suggest that subjects performing at this level could have saved a considerable amount of cognitive effort by simply choosing their actions at random.

Chapter Four presented a model that describes the problem-solving behavior observed in the laboratory experiments with human subjects. Highlights of the model include its simplicity, particularly when compared to existing learning models and the identification of a two-stage approach to problem solving. The two-stage model of behavior in novel task environments both extends and simplifies the existing theories of learning. The concept of mental models is discarded in favor of an explicit set of simple decision rules driven by performance feedback and simple performance comparisons. Stage One can be thought of as learning about the payoff characteristics of the environment. Stage Two is like problem solving in the more traditional sense. Both the search radius and search intensity are considerably smaller than Stage One levels as decision-makers work toward maximizing the payoffs. The shift from Stage One to Stage Two occurs after the decision-maker has surveyed the environment and attained a minimal level of understanding about the payoff structure. Interestingly, in environments with feedback error or a decreasing underlying trend, the transition to Stage Two never takes place.

Limitations

There are a number of limitations to this research. The limitations include those that are inherent in the task and methodology used to collect the data as well as limitations of the two-stage behavioral model. The two-stage model presented in Chapter Four is not intended to be a comprehensive behavioral model for problem solving in novel task environments. Any generalization of the findings to other tasks beyond those studied is risky. The task used in these experiments is an over-simplification of real world problems. In business, managers are often confronted with problems that have more than five variables and interactions among those variables. Additionally, it is rare that a task environment is truly novel. A manager usually has at least some knowledge of the market, product, and customers that is used to help guide his or her decisions. Also, groups rather than individuals make most real world business strategy decisions. Important problems are usually handled by several executives in order to consider multiple points of view and take advantage of a wider range of perspectives. Group decision-making also has the added effect of spreading the risk, in the event that there is an unfavorable outcome. Additionally, environments where the actions of a decision-maker have no impact on the environment are rare. For example, lowering prices will typically prompt a reaction by competitors, thus changing the competitive environment. Another limitation arises from the use of student subjects. Students are asked to perform unfamiliar tasks and are unmotivated beyond receiving additional course credit. In many business situations, a manger has his or her reputation, or in some cases their job, at stake when making key decisions.

Opportunities for Future Research

During the course of performing this research, numerous ideas emerged that are worthy of serious consideration but were beyond the scope of a single dissertation. These opportunities for future research can be arbitrarily classified into six groups including 1) new process indicators to identify the transition from Stage One to Stage Two, 2) new task environments, 3) production functions with dynamic trends that change endogenously, 4) multiple dependent variables, 5) evaluation of how players adapt to new environments, and 6) multi-player games. These will be reviewed in order in the following paragraphs.

The use of new process indicator variables will improve the understanding of when and why many people shift their problem solving strategies from a high search radius to a low search radius in the middle of the task. The goal is to deepen the understanding of the change from Stage One to Stage Two behavior. One technique is the addition of a timer in the Genesis task to measure the elapsed time between player actions. It is thought that the transition to Stage Two takes place after an extended period of thought and reflection by the decision-maker. A time delay between the two patterns of behavior would suggest that the decision-maker is taking additional time to make an assessment of their actions and performance up to that point in the task.

Another opportunity for future research is to identify and study interesting new task factors using essentially the same software and methodology already developed for this dissertation. These factors may include feedback delays and alternate forms of complexity. Feedback delays have been found to negatively affect human subjects' performance (Brehmer, 1991). It is not known whether any research has been done to evaluate feedback delays using simulated subjects. It is likely that some simulated decision strategies would be better suited than others to task environments with feedback delays. New forms of highly active strategies could be created to determine the best strategy for coping with delay. With human subjects, feedback delays could take at least two forms. One would be to delay the reporting of competitor performance data for one or more time periods. Another way to study feedback delay is to delay the reporting of the player's own performance data by some length of time.

A second task factor worthy of future research includes alternative forms of complexity. One approach is the addition of more player-controlled variables. The addition of more variables would increase the difficulty of the task compared to the current five-variable environment. The increased difficulty is expected to negatively impact performance for both simulated strategies and human subjects. Another technique for operationalizing complexity is the creation of environments where variables interact. The use of interactions among task variables would also have the effect of increasing the difficulty of the task. In many real world environments, variables interact with each other in various ways. For example, a retail manager dealing with an oversupply problem must reduce prices and increase advertising in order to move the merchandise. Doing either one without the other would only lead to a decrease in profits.

In the research reported here, the dynamic trend function changed exogenously and was not affected by the decision maker's actions. A potentially interesting variation of the task would be the use of a dynamic trend function that changes based on actions taken by the player. Relatively minor changes to the production function would be required to add this capability to the software. Future experiments could vary the degree to which player actions impact the environment and ultimately, performance. It is thought that learning would be differentially affected by the extent to which a decision-maker's actions affect the task environment.

Brehmer (1992) has studied the effects of multiple dependent variables in the context of the fire-fighting task. In addition to limiting the area lost to fire, subjects were also told to minimize the cost of fighting the fires. In the Genesis task, the profit objective could be supplemented with a goal to maximize market share. Players would be forced to make trade-offs in an effort to balance what may be perceived as conflicting objectives.

Another interesting avenue for future work is the evaluation of subjects' strategy changes from one run to the next in different environments. For example, subjects could play their first game in an environment where a low search radius leads to the best performance, i.e. the baseline environment. The second run would take place in an environment where the most successful strategy is to use a high search radius, such as the error environment. Some interesting questions can be posed. Under what conditions do people adapt their strategy to the environment? Assuming the player is successful in the first task using a systematic, low search radius approach, will they persist in the use of the same strategy in environments where it does not perform well? The results may have interesting implications relative to the research on the adaptive decision-maker (Payne, Bettman, & Johnson; 1993).

A number of interesting avenues for future research can be found in the literature on organization theory (OT). As noted in Chapter One, the five major streams of research in OT hint at how managerial decisions can affect the way organizations interact with the environment. The objective of this line of research would be to study the effects of environmental uncertainty, dynamism, and complexity in the context of the five theories. To properly evaluate the theories, a multi-player business strategy game will need to be designed and developed. Later, the results of the laboratory experiments could be validated with a field study to determine whether the same phenomena exist in the real world. Compared to the other potential avenues for future research, this is the most complicated and would require a fair amount of research and computer programming. However, it may also have the greatest potential to contribute to existing knowledge. If successful, the findings from the experiments and field study could add to the knowledge in both decision-making and organization theory.

APPENDIX ONE - DEFINITIONS OF TERMS

Action	The settings applied in a time period.
Complexity	Complexity can be operationalized in at least three different ways. One approach is to manipulate the number of task variables in the production function. A low complexity task would contain a small number of non-zero task variables, whereas a moderate to high complexity task would have two or three times as many non-zero coefficient variables. A second definition is the presence or absence of interactions among variables in the production function. The third way to operationalize complexity is through the use of non-linear coefficients in the production function.
Computational demand	A measure of the number of calculations required to execute each strategy.
Decision rule	An if-then conditional statement used by a strategy to determine the next action based on prior period performance.
Explorer	A highly active strategy characterized by large step sizes, a high degree of systematic search, variable levels of feedback aggregation, and moderate demands for memory and computation capability. After an initial evaluation of each variable, the explorer evolves into a hillclimber (or hill-descender) strategy.
Feedback accuracy	A measure of the amount of error contained in the outcome feedback. It is operationalized as the level of error in the production function.
Feedback aggregation (dwell period)	A feature of HASs, whereby settings in one period are re-applied in the following period(s). This enables payoffs to be averaged. A technique for coping with feedback error.
Forrest Gump	A name given to a highly active strategy characterized by large step sizes, no systematic search, no capability for feedback aggregation, and no demands for memory and computation capability. Uses a purely random combination of settings in each period.
Highly active strategy (HAS)	A set of decision rules that govern what action is to be taken in the next period(s) conditional upon prior period performance. HAS are a family of strategies based on varying degrees adaptive feedback and trial and error experimentation. At present, there are seven HAS: hillclimber, explorer, random-best-yet, Forrest Gump, and variants of the first three

	strategies that incorporate multiple period feedback aggregation.
Hillclimber	A highly active strategy characterized by small step sizes, a high degree of systematic search, variable levels of feedback aggregation, and moderate demands for memory and computation capability.
Intensity level	The setting applied to a variable in a particular period. A complete set of intensity levels settings constitutes an "action".
Memory requirement	The number of prior payoffs and settings that are required to be stored in memory in order to execute a specific HAS.
Performance	Performance can be measured in a number of ways. In the simplest sense, it is defined as the cumulative payoff at the end of period 30. Other measures of performance include: the payoff in the final period and the cumulative payoff in the early periods of the task.
Random-best-yet (RBY)	A highly active strategy characterized by large step sizes, a low degree of systematic search, variable levels of feedback aggregation, and low demands for memory and computation capability.
Search randomness	The level of logic used to determine the next action. Can be classified as either low (systematic) or high (random).
Step size	The magnitude of change in intensity level between periods for a single variable. For example, small step size strategies move a variable up or down one intensity level per period.
System dynamics	Exogenous change in the task environment that occurs over time, independent of player actions. Can be linear (increasing or decreasing) or a non-linear function. May be used to represent economic trends, the product life cycle, or other real world phenomenon.
Task environment	The production function.
Uncertainty	The decision-maker is unaware of future action-outcome linkages.

APPENDIX TWO - TABLES

Table 1

Dimensions of Highly Active Strategies

Dimension	Step Size	Search Randomness	Feedback Aggregation	No. of variables changed per period	Memory/ Computing Demands
Hillclimber	0 or 1	Low	None	1 or 2	Moderate
RBY	0 to 4	High	None	0 to 5	Low
Forrest	0 to 4	None	N/A	0 to 5	None
Explorer	1 or 4	Low	None	1 or 2	Moderate
Hillclimber 2	0 or 1	Low	2 period	0, 1, or 2	Moderate
RBY 2	0 to 4	High	2 period	0 to 5	Low
Explorer 2	0, 1, or 4	Low	2 period	0, 1, or 2	Moderate

Table 2

Experiment One: Baseline Environment

Strategy	Performance % of Maximum*	Ratio vs. Random
Forrest Gump (FG)	38% ± 1.5%	1.0
Hillclimber 2 (HC2)	43% ± 0.0%	1.1
Random-Best-Yet 2 (RBY2)	48% ± 3.1%	1.3
Random-Best-Yet (RBY)	51% ± 2.3%	1.3
Explorer 2 (EXP2)	60% ± 0.5%	1.6
Hillclimber (HC)	71% ± 0.0%	1.9
Explorer (EXP)	77% ± 0.8%	2.0

N = 350

* Confidence intervals calculated using $\alpha = .01$

Table 3

Experiment Two: Feedback Accuracy

	$s_e = 3$	$s_e = 6$	$s_e = 12$
Strategy	Performance % of Maximum*	Performance % of Maximum*	Performance % of Maximum*
Forrest Gump	38% ± 1.3%	38% ± 1.7%	39% ± 2.7%
Hillclimber 2	31% ± 3.3%	22% ± 3.0%	19% ± 4.5%
Random-Best-Yet 2	48% ± 3.4%	47% ± 3.8%	48% ± 3.9%
Random-Best-Yet	50% ± 2.4%	47% ± 2.8%	46% ± 3.5%
Explorer 2	58% ± 1.0%	57% ± 1.3%	38% ± 3.2%
Hillclimber	35% ± 4.7%	23% ± 4.9%	22% ± 4.3%
Explorer	68% ± 1.9%	67% ± 1.9%	45% ± 3.1%
	Ratio vs. Random	Ratio vs. Random	Ratio vs. Random
Forrest Gump	1.0	1.0	1.0
Hillclimber 2	0.8	0.6	0.5
Random-Best-Yet 2	1.3	1.2	1.3
Random-Best-Yet	1.3	1.2	1.2
Explorer 2	1.5	1.5	1.0
Hillclimber	0.9	0.6	0.6
Explorer	1.8	1.8	1.2

N = 1050

* Confidence intervals calculated using $\alpha = .01$

Table 4

Experiment Three: Complexity

Strategy	Function #1: Performance % of Maximum*	Function #1: Ratio vs. Random	Function #2: Performance % of Maximum*	Function #2: Ratio vs. Random
Forrest Gump (FG)	37% ± 1.3%	1.0	38% ± 1.4%	1.0
Hillclimber 2 (HC2)	49% ± 0.0%	1.3	39% ± 0.0%	1.0
Random-Best-Yet 2 (RBY2)	47% ± 2.8%	1.2	49% ± 3.2%	1.3
Random-Best-Yet (RBY)	51% ± 2.0%	1.3	50% ± 2.1%	1.3
Explorer 2 (EXP2)	51% ± 0.5%	1.3	56% ± 0.6%	1.5
Hillclimber (HC)	74% ± 0.0%	1.9	55% ± 0.0%	1.4
Explorer (EXP)	63% ± 0.8%	1.7	75% ± 0.9%	2.0

N = 700
* Confidence intervals calculated using $\alpha = .01$

Table 5

Experiment Four: System Dynamics - Increasing Linear Trend

Strategy	Slope = +2 Performance % of Maximum*	Slope = +2 Ratio vs. Random	Slope = +4 Performance % of Maximum*	Slope = +4 Ratio vs. Random
Forrest Gump (FG)	37% ± 1.5%	1.0	36% ± 1.3%	1.0
Hillclimber 2 (HC2)	19% ± 0.0%	0.5	19% ± 0.0%	0.5
Random-Best-Yet 2 (RBY2)	48% ± 3.0%	1.3	43% ± 3.4%	1.1
Random-Best-Yet (RBY)	51% ± 2.4%	1.3	48% ± 2.6%	1.3
Explorer 2 (EXP2)	50% ± 0.0%	1.3	50% ± 0.0%	1.3
Hillclimber (HC)	65% ± 2.5%	1.7	44% ± 0.0%	1.2
Explorer (EXP)	59% ± 1.6%	1.6	25% ± 0.0%	0.7

N = 700
* Confidence intervals calculated using $\alpha = .01$

Table 6

Experiment Four: System Dynamics - Decreasing Linear Trend

Strategy	Slope = -2		Slope = -4	
	Performance % of Maximum*	Ratio vs. Random	Performance % of Maximum*	Ratio vs. Random
Forrest Gump (FG)	37% ± 1.4%	1.0	38% ± 1.5%	1.0
Hillclimber 2 (HC2)	9% ± 0.0%	0.2	1% ± 0.0%	0.0
Random-Best-Yet 2 (RBY2)	44% ± 3.9%	1.2	41% ± 3.2%	1.1
Random-Best-Yet (RBY)	46% ± 3.0%	1.2	43% ± 3.5%	1.1
Explorer 2 (EXP2)	63% ± 0.0%	1.7	57% ± 0.0%	1.5
Hillclimber (HC)	39% ± 1.3%	1.0	9% ± 0.0%	0.2
Explorer (EXP)	80% ± 0.4%	2.1	73% ± 0.0%	1.9

N = 700

* Confidence intervals calculated using $\alpha = .01$

Table 7

Experiment Five: Means of Dependent Measures Across Blocks of Ten Time Periods

	Time Period			
	1 - 10	11 - 20	21 - 30	1 - 30
Performance (% of maximum)				
No error, no trend	45%	56%	61%	54%
Moderate error, no trend	45%	42%	49%	45%
No error, positive trend	41%	50%	52%	48%
Moderate error, positive trend	34%	40%	46%	40%
Grand mean	41%	47%	52%	47%
Avg. no. of variables changed per period				
No error, no trend	2.0	1.9	1.7	1.9
Moderate error, no trend	1.8	2.0	1.9	1.9
No error, positive trend	2.4	1.7	1.5	1.9
Moderate error, positive trend	2.7	1.8	1.3	1.9
Grand mean	2.2	1.8	1.6	1.9
Avg. change in intensity level per variable per period				
No error, no trend	1.3	1.3	1.5	1.3
Moderate error, no trend	1.2	1.3	1.3	1.3
No error, positive trend	1.5	1.3	1.3	1.4
Moderate error, positive trend	1.4	1.2	1.2	1.3
Grand mean	1.3	1.3	1.3	1.3
No. of periods with duplicate settings				
No error, no trend	2.3	2.5	4.0	11.7
Moderate error, no trend	2.9	2.3	3.2	11.4
No error, positive trend	1.3	2.9	3.6	9.8
Moderate error, positive trend	0.4	1.6	3.5	7.6
Grand mean	1.7	2.3	3.5	10.1
No. of periods with consecutive duplicates				
No error, no trend	1.8	1.6	2.9	6.3
Moderate error, no trend	1.9	1.3	2.1	5.6
No error, positive trend	0.9	2.2	2.7	6.1
Moderate error, positive trend	0.2	1.0	2.2	3.6
Grand mean	1.2	1.5	2.4	5.4

N = 20 for each group (10 subjects per group x 2 runs per subject)

Table 8

Experiment Five: Summary Of F Statistics From Univariate ANOVA Pooled Across All Time Periods And Simulation Runs

Factor	Performance (% of max)	Avg. no. of variables changed per period	Avg. change in intensity level per variable per period	No. of periods with duplicate settings	No. of periods with consecutive duplicates
Feedback Error	8.08***	0.04	1.74	1.00	2.00
Positive Trend	4.36**	0.01	0.13	4.81**	0.88
Error x Trend	0.07	0.00	0.05	0.54	0.58

** $p < .05$
*** $p < .01$

Table 9

Experiment Five: Summary of F Statistics From Univariate ANOVA By Time Block Pooled Across Simulation Runs

	Performance			Variables Changed			Δ Intensity Level			Duplicates			Consecutive Duplicates		
Block	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3
Error	1.25	9.29 ***	7.59 ***	0.00	0.33	0.00	0.81	0.64	1.84	0.06	2.40	0.75	0.63	3.18 *	1.31
Trend	6.14 **	1.13	3.67 *	12.9 ***	1.47	3.25 *	5.66 **	2.30	2.69	18.2 ***	0.10	0.02	13.4 ***	0.23	0.00
Inter-action	0.97	0.25	0.85	2.63	0.07	1.13	0.01	0.92	0.25	3.34 *	1.54	0.47	1.16	1.41	0.07

* $p < .10$
** $p < .05$
*** $p < .01$

Table 10

Experiment Five: Summary of Within Groups Paired Sample T-Tests Comparing Block One to Block Three

	Performance 1 - 3	Variables Changed 1 - 3	Δ Intensity Level 1 - 3	Duplicates Settings 1 - 3	Consecutive Duplicates 1 - 3
Baseline	4.52***	1.67	1.32	2.49**	1.66
Error	1.58	0.62	2.36**	0.57	0.56
Trend	2.97***	3.41***	1.62	3.66***	3.37***
Error & Trend	3.06***	8.28***	3.73***	6.30***	3.72***

** $p < .05$
*** $p < .01$

Table 11

Experiment Five: Trajectory Analysis

	Performance	Variables Changed	Change in Intensity Level	Duplicate Settings	Consecutive Duplicates
Baseline	UU = <u>13</u> UD = 4 DU = 3 DD = 0	UU = 0 UD = 6 DU = 7 DD = <u>7</u>	UU = 4 UD = 8 DU = <u>5</u> DD = 3	UU = <u>6</u> UD = 4 DU = 8 DD = 2	UU = 3 UD = 2 DU = <u>9</u> DD = 6
Error	UU = 2 UD = 6 DU = <u>11</u> DD = 1	UU = 4 UD = <u>7</u> DU = 3 DD = 6	UU = 1 UD = <u>10</u> DU = 3 DD = 3	UU = 1 UD = 4 DU = <u>10</u> DD = 5	UU = 2 UD = 2 DU = <u>8</u> DD = 8
Trend up	UU = <u>12</u> UD = 4 DU = 1 DD = 3	UU = 0 UD = 3 DU = 6 DD = <u>11</u>	UU = 3 UD = 3 DU = 6 DD = <u>8</u>	UU = <u>7</u> UD = 5 DU = 4 DD = 4	UU = <u>4</u> UD = 6 DU = 5 DD = 5
Error and trend up	UU = <u>7</u> UD = 7 DU = 6 DD = 0	UU = 0 UD = 1 DU = 4 DD = <u>15</u>	UU = 0 UD = 7 DU = 7 DD = <u>6</u>	UU = <u>7</u> UD = 4 DU = 7 DD = 2	UU = <u>4</u> UD = 4 DU = 6 DD = 6

Underlined values denote group mean trajectories.

Table 12

Experiment Five: Means of Dependent Measures Across Simulation Runs

	Simulation Run		Difference
	Run #1	Run #2	Run #1 - #2
Performance (% of maximum)			
No error, no trend	50%	53%	3%
Moderate error, no trend	42%	49%	7%
No error, positive trend	46%	49%	3%
Moderate error, positive trend	40%	40%	0%
Grand mean	45%	48%	3%
Avg. no. of variables changed per period			
No error, no trend	2.0	1.5	-0.5
Moderate error, no trend	2.0	1.8	-0.2
No error, positive trend	2.1	1.6	-0.5
Moderate error, positive trend	1.9	1.9	0.0
Grand mean	2.0	1.7	-0.3
Avg. change in intensity level per variable per period			
No error, no trend	1.3	1.2	-0.1
Moderate error, no trend	1.2	1.4	0.2
No error, positive trend	1.4	1.4	0.0
Moderate error, positive trend	1.3	1.3	0.0
Grand mean	1.3	1.3	0.0
No. of periods with duplicate settings			
No error, no trend	11.2	11.1	-0.1
Moderate error, no trend	11.0	11.8	0.8
No error, positive trend	6.8	12.8	6.0
Moderate error, positive trend	7.2	8.0	0.8
Grand mean	9.0	10.9	1.9
No. of periods with consecutive duplicates			
No error, no trend	5.5	6.5	1.0
Moderate error, no trend	5.3	5.9	0.6
No error, positive trend	4.5	7.7	3.2
Moderate error, positive trend	4.5	2.6	-1.9
Grand mean	4.9	5.6	0.7

N = 20 for each group

Table 13

Experiment Five: Correlation Matrix – All Groups

	Performance	Variables Changed	Δ Intensity Level	Duplicates	Consecutive Duplicates
Performance	1.000				
Variables Changed	-.469**	1.000			
Δ Intensity Level	.254*	.057	1.000		
Duplicates	.177	-.501**	-.149	1.000	
Consecutive Duplicates	.104	-.564**	-.086	.824**	1.000

* $p < .05$
** $p < .01$

Table 14

Experiment Five: Variance Inflation Factor Statistics

Process Indicator Variable	VIF Statistic
Variables Changed	1.476
Change in Intensity Level	1.027
Number of Duplicates	3.198
Number of Consecutive Duplicates	3.462

Table 15

Experiment Five: Regression Analysis Using Total Performance as the Dependent Variable

Group	Standardized Coefficients: Variables Changed	Δ Intensity Level	Duplicates	Consecutive Duplicates	Adj. R^2
All Groups	-0.59***	0.29***	0.29*	-0.44**	.32
Baseline	-0.90***	0.15	-0.03	-0.30	.48
Error	-0.79**	0.38*	0.23	-0.78**	.43
Trend	-0.81***	0.08	0.32	-0.58	.48
Error and Trend	-0.44*	0.43*	0.01	-0.10	.14

* $p < .10$
** $p < .05$
*** $p < .01$

Table 16

Experiment Six: Means of Dependent Measures Across Blocks of Ten Time Periods

	Time Period			
	1 - 10	11 - 20	21 - 30	1 - 30
Performance (% of maximum)				
No trend, no complexity	53%	67%	71%	64%
Negative trend, no complexity	50%	54%	54%	53%
No trend, moderate complexity	48%	58%	64%	57%
Negative trend, moderate complexity	51%	50%	54%	52%
Grand mean	50%	58%	61%	56%
Avg. no. of variables changed per period				
No trend, no complexity	1.8	1.1	1.1	1.3
Negative trend, no complexity	1.9	2.2	2.4	2.2
No trend, moderate complexity	2.0	1.6	1.4	1.7
Negative trend, moderate complexity	2.2	2.3	2.4	2.3
Grand mean	2.0	1.8	1.8	1.9
Avg. change in intensity level per variable per period				
No trend, no complexity	1.5	1.6	1.4	1.6
Negative trend, no complexity	1.3	1.3	1.6	1.4
No trend, moderate complexity	1.1	1.1	1.1	1.1
Negative trend, moderate complexity	1.2	1.2	1.2	1.2
Grand mean	1.3	1.3	1.3	1.3
No. of periods with duplicate settings				
No trend, no complexity	3.4	4.8	5.7	16.4
Negative trend, no complexity	2.2	1.8	2.3	11.0
No trend, moderate complexity	2.3	3.9	4.5	13.8
Negative trend, moderate complexity	1.6	1.2	1.6	8.5
Grand mean	2.4	2.9	3.6	12.5
No. of periods with consecutive duplicates				
No trend, no complexity	2.4	3.6	4.4	10.8
Negative trend, no complexity	0.9	0.6	0.4	2.0
No trend, moderate complexity	1.2	2.6	3.4	7.6
Negative trend, moderate complexity	0.5	0.5	0.4	1.5
Grand mean	1.3	1.8	2.2	5.5

N = 90

Table 17

Experiment Six: Summary Of F Statistics From Univariate ANOVA Pooled Across All Time Periods And Simulation Runs

Factor	Performance (% of max)	Avg. no. of variables changed per period	Avg. change in intensity level per variable per period	No. of periods with duplicate settings	No. of periods with consecutive duplicates
Negative Trend	11.77***	25.54***	0.59	26.15***	61.26***
Complexity	2.50	2.68*	12.08***	5.78**	3.14*
Trend x Complexity	1.62	0.47	1.77	0.00	1.70

* $p < .10$
** $p < .05$
*** $p < .01$

Table 18

Experiment Six: Summary of F Statistics From Univariate ANOVA By Time Block Pooled Across Simulation Runs

	Performance			Variables Changed			Δ Intensity Level			Duplicates			Consecutive Duplicates		
Block	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3
Trend	0.00	16.1 ***	21.2 ***	0.57	29.0 ***	40.2 ***	0.01	0.64	0.91	5.77 **	41.8 ***	54.4 ***	8.58 ***	34.1 ***	68.6 ***
Complexity	0.37	6.16 **	1.45	1.58	2.93 *	1.11	6.02 **	9.31 ***	9.77 ***	4.51 **	2.98 *	4.76 **	4.42 **	1.51	0.92
Interaction	1.43	0.94	1.41	0.02	1.37	0.37	2.09	3.50 *	0.05	0.33	0.16	0.35	1.20	0.86	0.92

* $p < .10$
** $p < .05$
*** $p < .01$

Table 19

Experiment Six: Summary of Within Groups Paired Sample T-Tests Comparing Block One to Block Three

	Performance 1 - 3	Variables Changed 1 - 3	Δ Intensity Level 1 - 3	Duplicates Settings 1 - 3	Consecutive Duplicates 1 - 3
Baseline	4.71***	3.30***	0.28	3.04***	2.27**
Trend	1.74*	2.27**	2.68**	0.18	1.98*
Complexity	5.80***	4.88***	0.21	3.85***	3.67***
Trend & Complexity	1.28	1.99*	0.45	0.00	0.77

* $p < .10$
** $p < .05$
*** $p < .01$

Table 20

Experiment Six: Trajectory Analysis

	Performance	Variables Changed	Change in Intensity Level	Duplicate Settings	Consecutive Duplicates
Baseline	UU = 13 UD = 8 DU = 0 DD = 1	UU = 1 UD = 4 DU = 7 DD = 10	UU = 0 UD = 11 DU = 5 DD = 6	UU = 10 UD = 7 DU = 3 DD = 2	UU = 7 UD = 5 DU = 6 DD = 4
Trend down	UU = 9 UD = 6 DU = 3 DD = 4	UU = 7 UD = 7 DU = 4 DD = 4	UU = 8 UD = 5 DU = 8 DD = 1	UU = 4 UD = 1 DU = 9 DD = 8	UU = 0 UD = 2 DU = 3 DD = 17
Complexity	UU = 14 UD = 6 DU = 4 DD = 0	UU = 1 UD = 9 DU = 7 DD = 7	UU = 4 UD = 7 DU = 7 DD = 6	UU = 6 UD = 6 DU = 8 DD = 4	UU = 7 UD = 4 DU = 8 DD = 5
Trend down complexity	UU = 5 UD = 4 DU = 11 DD = 2	UU = 4 UD = 12 DU = 4 DD = 2	UU = 2 UD = 10 DU = 7 DD = 3	UU = 1 UD = 4 DU = 9 DD = 8	UU = 0 UD = 2 DU = 3 DD = 17

Underlined values denote group mean trajectories.

Table 21

Experiment Six: Means of Dependent Measures Across Simulation Runs

	Simulation Run		Difference
	Run #1	Run #2	Run #1 - #2
Performance (% of maximum)			
No trend, no complexity	62%	66%	4%
Negative trend, no complexity	54%	52%	-2%
No trend, moderate complexity	58%	55%	-3%
Negative trend, moderate complexity	53%	51%	-2%
Grand mean	57%	56%	-1%
Avg. no. of variables changed per period			
No trend, no complexity	1.4	1.2	-0.2
Negative trend, no complexity	2.4	2.0	-0.4
No trend, moderate complexity	1.7	1.6	-0.1
Negative trend, moderate complexity	2.6	2.1	-0.5
Grand mean	2.0	1.7	-0.3
Avg. change in intensity level per variable per period			
No trend, no complexity	1.5	1.6	0.1
Negative trend, no complexity	1.4	1.4	0.0
No trend, moderate complexity	1.2	1.1	-0.1
Negative trend, moderate complexity	1.2	1.3	0.1
Grand mean	1.3	1.4	0.1
No. of periods with duplicate settings			
No trend, no complexity	15.1	17.6	2.5
Negative trend, no complexity	10.6	11.4	0.8
No trend, moderate complexity	12.8	14.8	2.0
Negative trend, moderate complexity	7.7	9.3	1.6
Grand mean	11.6	13.3	1.7
No. of periods with consecutive duplicates			
No trend, no complexity	8.5	13.2	4.7
Negative trend, no complexity	1.2	2.7	1.5
No trend, moderate complexity	7.4	8.3	0.9
Negative trend, moderate complexity	0.5	2.5	2.0
Grand mean	4.4	6.7	2.3

N = 90

Table 22

Experiment Six: Correlation Matrix – All Groups

	Performance	Variables Changed	Δ Intensity Level	Duplicates	Consecutive Duplicates
Performance	1.000				
Variables Changed	-.571**	1.000			
Δ Intensity Level	.156	.015	1.000		
Duplicates	.478**	-.747**	.209*	1.000	
Consecutive Duplicates	.388**	-.722**	.139	.810**	1.000

* $p < .05$
** $p < .01$

Table 23

Experiment Six: Variance Inflation Factor Statistics

Process Indicator Variable	VIF Statistic
Variables Changed	2.666
Change in Intensity Level	1.125
Number of Duplicates	3.704
Number of Consecutive Duplicates	3.192

Table 24

Experiment Six: Regression Analysis Using Total Performance as the Dependent Variable

	Standardized Coefficients				
Group	Variables Changed	Δ Intensity Level	Duplicates	Consecutive Duplicates	Adj. R^2
All Groups	-0.59***	0.16*	0.15	-0.19	.34
Baseline	-0.63**	-0.13	0.26	-0.29	.33
Trend	-0.55*	0.37*	0.00	-0.41**	.43
Complexity	-0.79*	0.16	-0.40	0.22	.26
Trend and Complexity	0.12	-0.15	0.69**	0.05	.23

* $p < .10$
** $p < .05$
*** $p < .01$

Table 25

Comparison of Subjects' Performance To The Random Strategy

Experiment Five

Group	Performance (% of Maximum)			Percent Performing Better than Random*
	Lowest	Median	Highest	
No error, no trend	34%	54%	83%	85%
Moderate error, no trend	24%	45%	72%	70%
No error, positive trend	25%	45%	71%	70%
Moderate error, positive trend	21%	41%	56%	65%

N = 80
* Random strategy performs at 38% ± 1.5% ($\alpha = .01$) in all task environments

Experiment Six

Group	Performance (% of Maximum)			Percent Performing Better than Random*
	Lowest	Median	Highest	
No trend, no complexity	46%	64%	84%	100%
Negative trend, no complexity	26%	50%	77%	86%
No trend, moderate complexity	44%	57%	81%	100%
Negative trend, moderate complexity	31%	50%	77%	96%

N = 90

APPENDIX THREE - FIGURES

Figure 1

Genesis Business Strategy Simulation Software

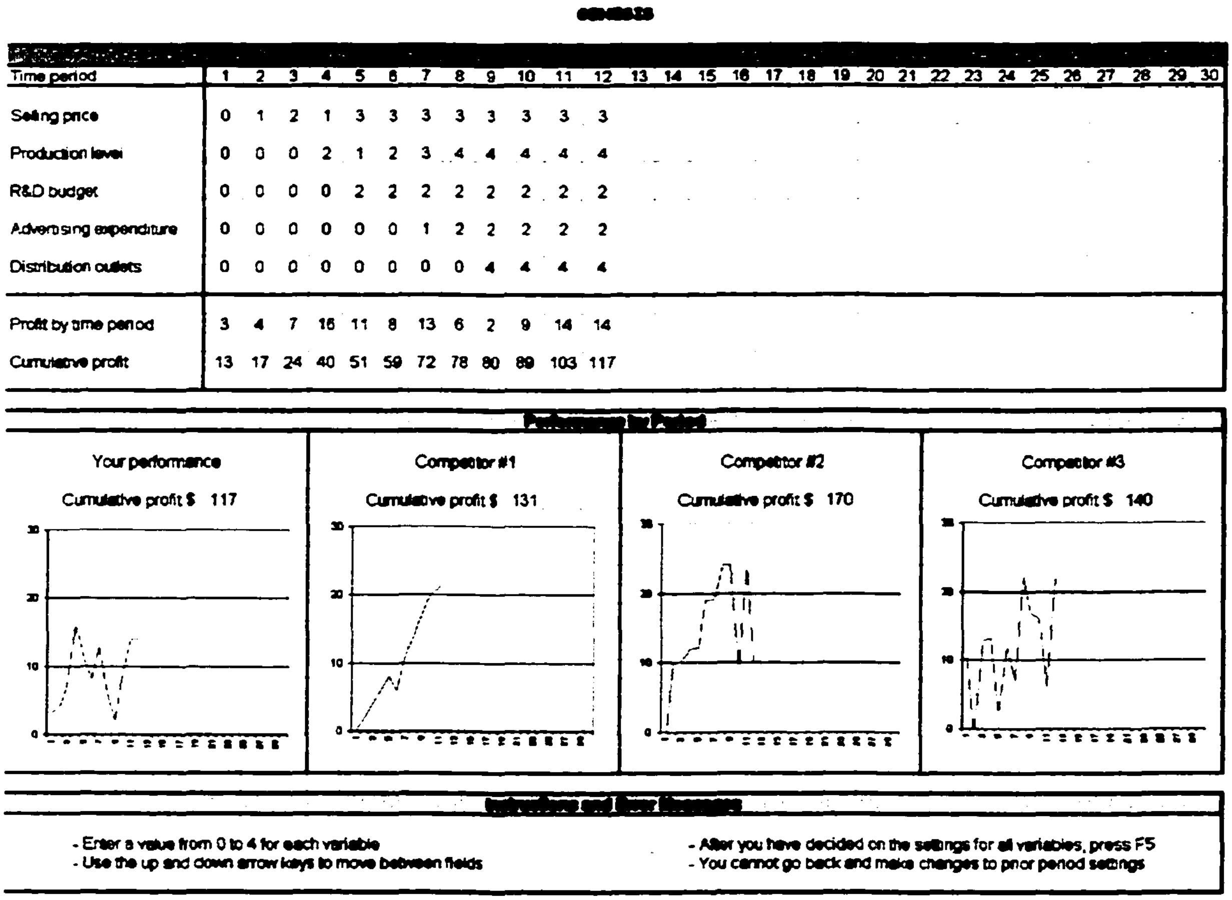

Time period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
Selling price	0	1	2	1	3	3	3	3	3	3	3	3																		
Production level	0	0	0	2	1	2	3	4	4	4	4	4																		
R&D budget	0	0	0	0	2	2	2	2	2	2	2	2																		
Advertising expenditure	0	0	0	0	0	0	1	2	2	2	2	2																		
Distribution outlets	0	0	0	0	0	0	0	0	4	4	4	4																		
Profit by time period	3	4	7	16	11	8	13	6	2	9	14	14																		
Cumulative profit	13	17	24	40	51	59	72	78	80	89	103	117																		

Figure 2

Feedback Loop Model
Forrester (1961)

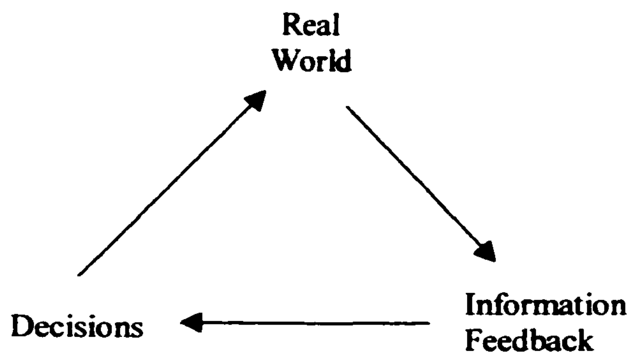

Figure 3

Single Loop Learning Model
Argyris (1985)

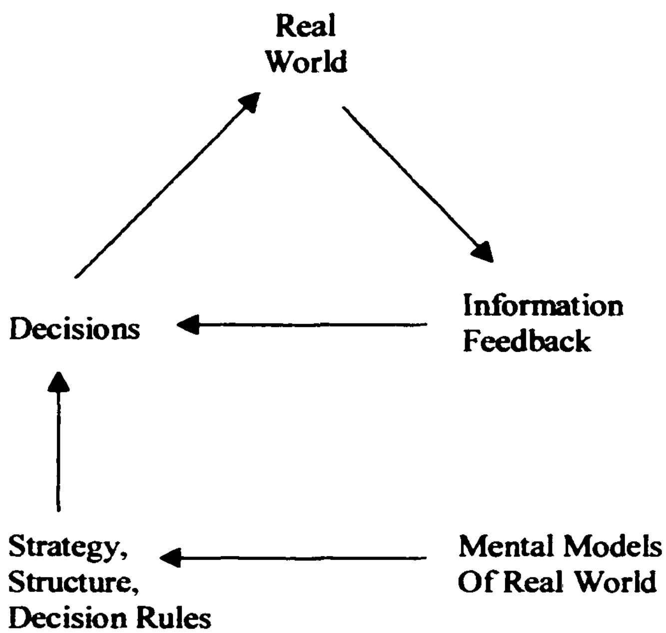

Figure 4

Double Loop Learning Model
Argyris (1985)

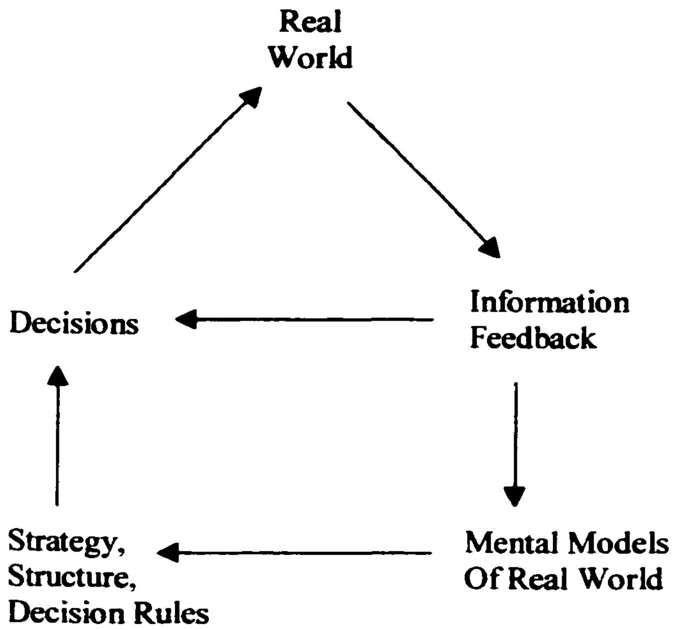

Figure 5

Two-Stage Model of Problem Solving Behavior in Novel Task Environments
Stage One

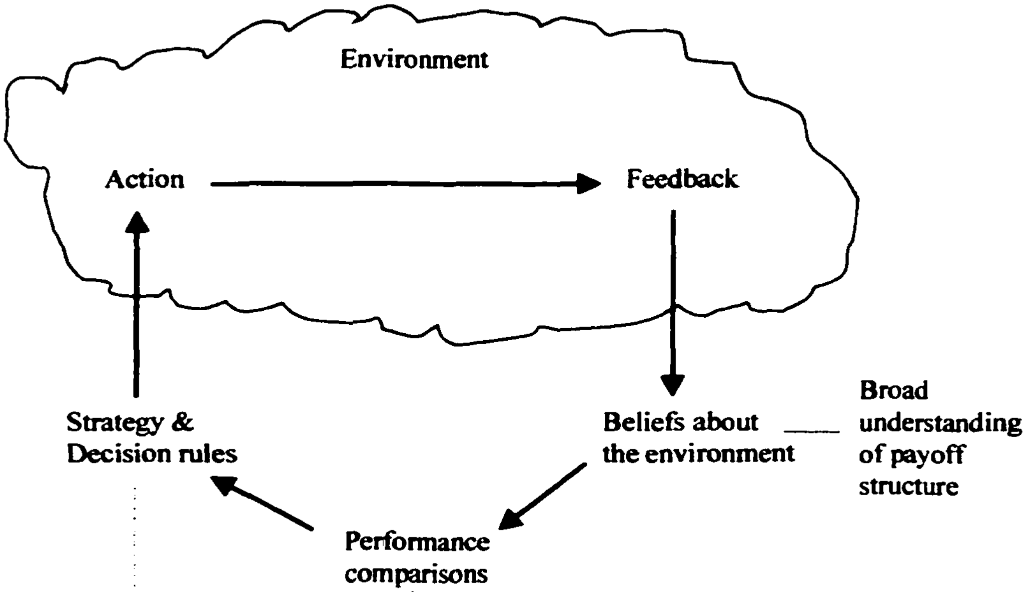

Slope of recent performance - increasing, decreasing, or zero
Sign of recent performance - positive or negative
Stability - yes or no (do duplicate settings produce identical profits?)
Linearity - yes or no (do increases/decreases in intensity level lead to proportionate changes in performance?)

<u>Search Radius Decision</u>
Is recent performance producing losses? If yes, increase search radius and go to next rule. If no, go to next rule.

Is recent performance declining? If yes, increase search radius and take next action. If no, go to next rule.

Is the environment stable? If yes, decrease search radius and advance to stage two. If no, maintain current search radius and take next action.

<u>Search Intensity Decision</u>
Are one or more variables non-linear? If yes, decrease search intensity and take next action. If no, maintain current search intensity and take next action.

Figure 6

Two-Stage Model of Problem Solving Behavior in Novel Task Environments
Stage Two

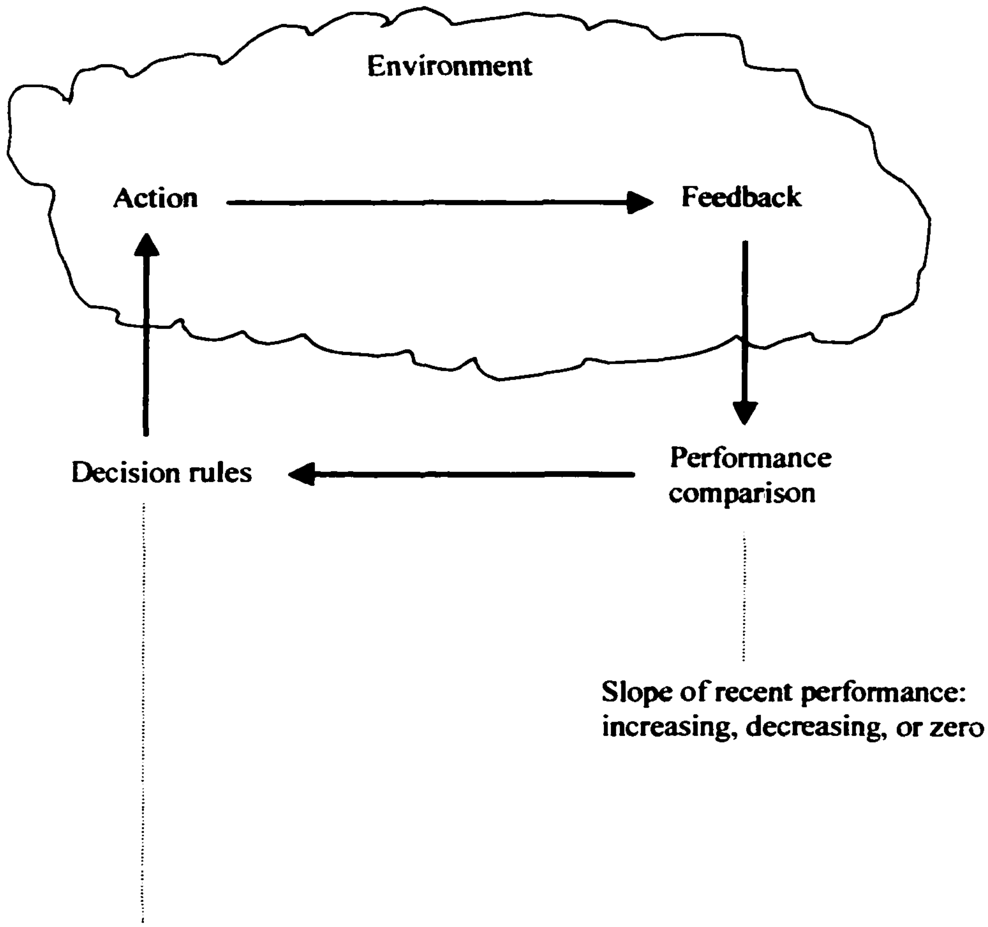

Is performance improving?

If yes, increment the active variable by one intensity level. Take next action.

If no, revert to prior settings that produced the highest payoff, select a new active variable, and increment the new active variable by one intensity level. Take next action.

REFERENCES

Aldrich, H. (1979). *Organizations and environments.* Prentice-Hall. Englewood cliffs, N.J.

Argyris, C. (1985). *Strategy, change, and defensive routines*. Boston: Pitman.

Brehmer, B. (1990). Strategies in real-time, dynamic decision making. In R. Hogarth (Ed.), *Insights in decision making. A tribute to Hillel J. Einhorn* (pp. 262-279). Chicago, IL: University of Chicago Press.

Brehmer, B. & Allard, R. (1991). Real-time dynamic decision making: Effects of task complexity and feedback delays. In J. Rasmussen, B. Brehmer, and J. Leplat (Eds.) *Distributed decision making, Cognitive models for cooperative work.* Chichester: Wiley.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, 81, 211-241.

Campbell, D. (1969). Reforms as experiments. *American Psychologist*. 409-429.

Child, J. (1972). Organizational structure, environment and performance: The role of strategic choice. *Sociology*, 6, 1-30.

Connolly, T. & Miklausich, V. (1978). Some effects of feedback error in diagnostic decision tasks. *Academy of Management Journal* 21: 301-307.

Connolly, T. & W. G. Wagner (1988). "Decision cycles." *Advances in Information-processing in Organizations* 3: 183-205.

Connolly, T. (1988). Hedge-clipping, tree-felling, and the management of ambiguity: The need for new images of decision-making. *Managing Ambiguity and Change*. L. R. Pondy, R. J. Boland, & T. Thomas. New York, Wiley: 37-50.

Connolly, T. (1999). Action as a fast and frugal heuristic. *Minds and Machines*. 9, 479-496.

Davis, G. F. & Powell, W. (1992). Organization-environment relations. In M. Dunnette (Ed.) *Handbook of Industrial and Organizational Psychology*. Vol. 3, 2nd ed., pp. 798-1033. Palo Alto, CA: Consulting Psychologists Press.

Dess, G. G., & Beard, D. W. (1984). Dimensions of organizational task environments. *Administrative Science Quarterly*, 29, 52-73.

Diehl, E. & Sterman, J. D. (1993). Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes*, 62, 198-215.

Dorner, D. (1980). On the problems people have in dealing with complexity. *Simulation and Games*, 11, 87-106.

Dorner, D. & Wearing, A.J. (1995). Complex problem solving: Towards a (computer-simulated) theory. In P.A. Frensch & J. Funke, (Eds.) *Complex problem solving: The European perspective* (pp. 65-99). Hillsdale, N.J.: Erlbaum.

Duncan, R. B. (1972). Characteristics of organizational environments and perceived environmental uncertainty. *Administrative Science Quarterly*, 17(3), 313-327.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors*, 4, 59-73.

Eisenhardt, K. M., & Bourgeios, L. J. (1988). Politics of strategic decision making in high-velocity environments: Toward a midrange theory. *Academy of Management Journal*, 31(4), 737-770.

Eisenhardt, K. M. (1989). Making fast strategic decisions in high-velocity environments. *Academy of Management Journal*, 32, 543-578.

Eisenhardt, K. M., & Zbaracki, M. (1992). Strategic decision making. *Strategic Management Journal*, 13, 17-37.

Emery, F. E., & Trist, E. L. (1965). The causal texture of organizational environments. *Human Relations*, 18, 21-32.

Forrester, J. W. (1961). *Industrial dynamics*. Portland, OR: Productivity Press.

Forrester, J. W. (1992). Policies, decisions, and information sources for modeling. *European Journal of Operational Research*, 59: 42-63.

Frensch, P.A. and Funke, J. (Eds.). (1995). *Complex problem solving: The European perspective*. Hillsdale, N.J.: Erlbaum.

Funke, J. & Müller, (1988). Active control and prediction as determinants of system identification and system control. *Sprache and Kognition*, 7, 176-186.

Funke, J. (1991). Solving complex problems: Exploration and control of complex systems. In R. J. Stemberg and P. A. Frensch (Eds.) *Complex problem solving: Principles and mechanisms*. Hillside, N.J.: Erlbaum.

Funke, J. (1992). Dealing with dynamic systems: Research strategy, diagnostic approach and experimental results. *The German Journal of Psychology*, 16, 24 - 43.

Funke, J. (1993). Microworlds based on linear equation systems: a new approach to complex problem solving and experimental results. In G. Strube and K.-f. Wender, eds. *The cognitive psychology knowledge. The German Wissenpsychologie Project.* Amsterdam. Elsevier Science Publications.

Funke, J. (1995). Experimental research on complex problem solving. *Complex problem solving: The European perspective*. Hillsdale, N.J.: Erlbaum.

Gibson, F., Fichman, M. & Plaut, D. C. (1997). Learning in dynamic decision tasks: Computational model and empirical evidence. *Organizational Behavior and Human Performance,* 71, 1-35.

Gigerenzer, G. and Goldstein, D.D. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650-660.

Gigerenzer. G. & Todd, P. (1999). *Simple Heuristics That Make us Smart.* Oxford, England. Oxford University Press.

Hannan, M. & Freeman, J. (1977). The population ecology of organizations. *American Journal of Sociology*, 82, 929-966.

Heineken, E., Arnold, H.-J., Kopp, A., & Soltysiak, R. (1992). Strategies of thinking in controlling a simple dynamic system under different deadtime conditions. *Sprache & Kognition*, 11, 136-148.

Hesse, F. W. (1982). Effects of semantic context on problem solving. *Zeitschrift fur Experimentelle und Angewandte Psychologie, 29, 62-91.*

Hogarth, R. & Makridakis, S. (1981). The value of decision making in a complex environment: An experimental approach. *Management Science.* 27:1, 93-107.

Huber, 0. (1994a). Decision behavior in a multistage investment task. *Acta Psychologica*, 85, 139-154.

Huber, 0. (1994b). Simulation as a tool to improve process models. *Acta Psychologica*, 87, 121 - 134.

Huber, 0. (1995). Complex problem solving as multistage decision making. In P.A. Frensch & J. Funke, (Eds.) *Complex problem solving: The European perspective* (pp. 151 - 173). Hillsdale, N.J.: Erlbaum.

Huber, 0 (1996). Buying protection in a multistage investment task. *Acta Psychologica*, 92, 153-167.

Huber, O., Debeutz, A, Pratscher, J. and Quehenberger, I. (1990). Perceived Control in a Multistage Decision Task. *Journal of Behavioral Decision Making*, 3, 123-136.

Kerstholt, J. H. (1996). The effects of information costs on strategy selection in dynamic tasks. *Acta Psychologica*, 94, 273 -290.

Kerstholt, J. H., & Raaijmakers, J. G. W. (1997). Decision making in dynamic task environments. In W. R. Crozier & 0. Svenson (Eds.) *Decision making: Cognitive models and explanations*. London: Routledge. (Pp. 205-217).

Kleinmuntz, D. & Kleinmuntz, B. (1981). Systems simulation decision strategies in simulated environments. *Behavioral Science*, 26, 255-270.

Kleinmuntz, D. & Thomas, J. (1987). The value of action and inference in dynamic decision making. *Organizational Behavior and Human Decision Processes*, 39, 341-364.

Kleinmuntz, D. (1985). Cognitive heuristics and feedback in a dynamic decision environment. *Management Science*, 31, 680-702.

Kleiter, G. D. (1975) Dynamic decision behavior: Comments on Rapoport's paper. In D. Wendt & C. Vlek (Eds.) *Utility, Probability, and Human Decision Making*, Dordrecht-Holland: Reidet. (Pp. 371-380).

Köller, O., Dauenheimer D. G., & Strauß, B. (1993). Differences between individuals and two-person groups in solving complex problems in relation to initial ability level. *Zeitschrift für Experimentelle und Angewandte Psychologie*, 40:194-221.

Lawrence, P. R. & Lorsch, J. W. (1967). *Organization and Environment*. Boston: Harvard University Press.

Leutner, D. (1988). Computer simulated dynamic systems: Knowledge acquisition under different forms of teaching and types and social context. *Zeitschrift fur Entwicklungpsychologie und Padagogische Psychologie*, 40:338-355.

Lindblom, C. E. (1979). Still muddling, not yet through. *Public Administration Review*, 39:517-526.

Luenberger, D. G. (1979). *Introduction to dynamic systems*. N.Y.: Wiley.

Mackinnon, A. J. & Wearing, A. J. (1980). Complexity and decision making. *Behavioral Science*, 25, 285-296.

MacKinnon, A.J. and Wearing, A.J. (1985). Systems analysis and dynamic decision making. *Acta Psychologica*, 58, 159-172.

Miles, R., Snow, C, & Pfeffer, J. (1974). Organization-environment: Concepts and issues. *Industrial Relations*, 13:244-264.

Miller, D. (1987b). Strategy making and structure: Analysis and implications for performance. *Academy of Management Journal*, 30(1), 280-308.

Mitroff, I. & Mason, R. (1981). *Creating a dialectical social science: Concepts, methods, and models*. Dordrecht, Holland. Reidel.

Nelson, R. & Winter, S. (1982). *An evolutionary theory of economic change*. Cambridge, MA. Harvard University Press.

Neter, J., Kutner, M., Nachtsheim, C., & Wasserman, M. (1996). *Applied Linear Statistical Models*. Chicago, IL. McGraw-Hill.

Paich, M. & Sterman, J. D. (1993). Boom, bust, and failures to learn in experimental markets. *Management Science*, 39 (12): 1439:1458.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. NY: Cambridge University Press.

Perrow, C. (1972). Complex organizations: A critical essay. New York: McGraw-Hill.

Pfeffer, J. & Salancik, G. (1978). *The external control of organizations: A resource dependence perspective*. New York: Harper and Row.

Powell, W. & DiMaggio, P. (Eds.). (1991). *The new institutionalism in organizational analysis*. Chicago: University of Chicago Press.

Priem, R. L. (1994). Executive judgment, organizational congruence, and firm performance. *Organizational Science*. (5) 421-437.

Putz-Osterloh, W. (1981). On the relationship between test intelligence and success in problem solving. *Zeitschrift für Psychologie*, 189, 79-100

Putz-Osterloh, W., & Lüer, G. (1981). On whether results from a test of intelligence can predict problem solving performance. *Zeitschrift für Experimentelle und Angewandte Psychologie*, 28, 309-334.

Quinn, J. B. (1989). Managing strategies incrementally. In H. Leavitt, Pondy, L., and

Boje, D. (eds). *Readings in Managerial Psychology*, Chicago: University of Chicago Press.

Rapoport, A. (1967). Dynamic programming models for multistage decision making. *Journal of Mathematical Psychology*, 4, 48-71.

Rapoport, A. (1975). Research paradigms for studying dynamic decision behavior. In D. Wendt & C. Vlek (Eds.) *Utility, Probability, and Human Decision Making*, Dordrecht-Holland: Reidel. (Pp. 347-369).

Reger, R. & Palmer, T. (1996). Managerial categorization of competitors: Using old maps to navigate new environments. *Organization Science*. 7:22-39.

Reither, F. (1981). About thinking and acting of experts in complex situations. *Simulation & Games*, 12, 125-140.

Rost, J., & Strauß, B. (1993). Interaction of information feedback and mental model in solving a complex problem. *Sprache & Kognition, 12, 73--82.*

Starbuck, W. (1976). Organizations and their environments. In Marvin Dunnette (ed.), *Handbook of Industrial and Organizational Psychology*: 1069-1123. Chicago; Rand-McNally.

Sterman, J. D. (1987). Testing behavioral simulation models by direct experiment. *Management Science*, 33, 12:1572-1592.

Sterman, J. D. (1989c). Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes*, 43, 301-335.

Sterman, J. D. (1994). Learning in and about complex systems. *System Dynamics Review*, 10, 291-330.

Strohschneider, S. (1990). *Knowledge acquisition and action regulation*. Wiesbaden, Germany: Deutscher Universitats Verlag.

Strohschneider, S. (1991). No system of systems! Reply to the paper "System features as determinants of behavior in dynamic task environments" by Joachim Funke. *Sprache & Kognition*, 10, 109-113.

Süß, H.-M., Oberauer, K., & Kersting, M. (1993). Intelligence and control performance on computer-simulated systems. *Sprache & Kognition*, 12, 83-97.

Svenson, 0. (1996). Decision making and the search for fundamental psychological regularities: What can be learned from a process perspective? *Organizational Behavior* *and Human Decision Processes*, 65, 252-267.

Terreberry, S. (1968). The evolution of organizational environments. *Administrative Science Quarterly*, 12:377-396.

Thompson, J. (1967). *Organizations in Action*. New York: McGraw Hill.

Toda, M. (1962). The design of the fungus eater: A model of human behavior in an unsophisticated environment. *Behavioral Science*, 7, 164-183.

Venkatraman, N., & Prescott, J. E. (1990). Environment-strategy coalignment: An empirical test of its performance implications. *Strategic Management Journal*, 11(1), 1-24.

Vollmeyer, R., & Holyoak, K. J. (1993). Hypothesis testing strategies in learning a complex dynamic system. Paper presented at the Fifth Annual Convention of the American Psychological Society, Chicago, IL.

Williamson, O. (1975). *Markets and hierarchies*. New York: Free Press.

Williamson, O. (1981). The economics of organizations: The transaction cost approach. *American Journal of Sociology*, 87, 548-577.

You, T. (1989). Disclosing the decision-maker's internal model and control policy in a dynamic decision task using a system control paradigm. Unpublished MA thesis, Purdue University.