BOOM AND BUST: DECISION MAKING IN A DYNAMIC MARKET ENVIRONMENT

by

Mark Paich

M.A., University of Colorado
(1976)

SUBMITTED TO THE SLOAN SCHOOL OF MANAGEMENT
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

December 1994

© 1994 Mark Paich
All Rights Reserved

The author hereby grants to M.I.T. permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part.

Signature of Author _____
Sloan School of Management
December, 1994

Certified by _____
John Sterman
Thesis Supervisor

Accepted by _____
Chair, Departmental Committee on Graduate Studies

ARCHIVES
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
JAN 30 1995

BOOM AND BUST: DECISION MAKING IN A DYNAMIC MARKET ENVIRONMENT

Mark Paich

Submitted to the Department of Management
In December 1994 in partial fulfillment of the
Requirements for the Degree of Doctor of Philosophy

ABSTRACT

Economists and psychologists have both tested the rationality of decision makers in numerous experiments. In behavioral decision theory, the vast majority of these experiments have involved one-time choices in static environments. Recently, it has been questioned whether experimental results from static environments can be generalized to the dynamic environments that characterize many real decision tasks. In dynamic situations, decision makers make a series of choices over time based on feedback about the outcome of previous decisions.

Recently, a number of experiments in behavioral decision theory have dealt with dynamic environments. Several of these experiments have shown that decision makers perform very poorly in complex dynamic environments, especially when delays and feedback relationships are present.

Generally, these experimental studies of dynamic decision making have several limitations. First, the causal structure of the task environment is very simple. Relatively few experiments include the delays and closed feedback processes that characterize real systems. Second, feedback complexity, time delays and non-linearity are not varied as experimental treatments which makes it difficult to attribute performance changes to the structure of the experimental environment. Finally, the effect of learning is not explicitly measured. Several of the experiments involve only single trials and do not allow for the possibility that, with experience, subjects learn improved decision rules.

This thesis presents an experiment that deals directly with these three limitations. In the experiment, subjects manage a simulated firm's investment and pricing decisions over the course of a product's life cycle. The market simulation used in the experiment is based on a well-established and realistic model of the product life-cycle that includes many of the feedback structures found in real markets. By varying parameters in the model, it is possible to evaluate how decision makers' performance changes with the strength of important feedback loops.

Thesis Supervisor: Dr. John Sterman

Title: Professor of Management Science

Acknowledgments

Many thanks to Becky, Julia and Ryan for their support and patience during a long and difficult process. A special thanks to my parents Ralph and Daisy for their encouragement, support and total confidence that this story could have a happy ending. For all of you, this is as much your achievement as it is mine.

I am indebted to John Sterman in so many ways that I can't begin to list them. Without John's ideas, insightful criticisms, and encouragement, there is no chance that this thesis would have been completed. I cannot imagine a better advisor.

I also wish to thank John Carroll and Michael Scott Morton for the time they spent reading and critiquing the work. Your comments were very helpful in forcing me to explain why these ideas might be significant.

Thanks to Rod for his work in programming the simulation interface and in handling the data.

Finally, thanks to Barry Richmond for introducing me to system dynamics and for being part of a true intellectual partnership. Thanks to George Richardson and Nelson Repenning for your support on the day of the defense. Thanks to Roy Hinton and Norm Torre for their support in the last days of this project.

CONTENTS

	Page
Abstract	2
Acknowledgments	3
Table of Contents	4
1. Introduction	5
2. Review of the Literature	11
3. Case Studies in Growth Management	33
4. Simulation Model	55
5. Experimental Design	78
6. Discussion of Results	109
References	152

Chapter One
Introduction

Economists and psychologists have long debated the question of rationality in decision making (Hogarth & Reder 1986, Simon 1982). Economists view decision makers as rational actors who have clear preferences and efficiently process information into optimal choices (Simon 1986). In the extreme, economic models assume that decision makers possess accurate mental models of extremely complex dynamic systems such as the macro economy (Sargent 1979).

In contrast, many psychologists believe that decision makers can be characterized by bounded rationality (Simon 1982). Decision makers faced with even moderately complex situations resort to simplifying heuristics which systematically produce sub-optimal outcomes (Hogarth & Makridakis 1981). Research in behavioral decision theory has identified a long list of heuristics used in commonly encountered choice situations.

Economists and psychologists have both tested the rationality of decision makers (DM) in numerous experiments. In behavioral decision theory, the vast majority of these experiments have involved one-time choices in static environments (Hogarth 1981, Brehmer 1987). In a 1981 review paper, Hogarth questioned whether experimental results from static environments can be generalized to the dynamic environments that characterize many real decision tasks. In dynamic situations, DMs make a series of choices over time based on feedback about the outcome of previous decisions. As Hogarth states,

> "With few exceptions, judgment researchers have focused on discrete incidents (particular actions, predictions and choices) that punctuate these continuous processes; furthermore, task environments are typically conceptualized to be stable. Although judgmental performance has been evaluated according to the principles of optimal behavior implied by decision theory and the probability calculus, few researchers have questioned whether the

> assumptions of such models apply to continuous processes." (Hogarth 1981)

Research in behavioral decision theory has not yet determined whether DMs perform better in dynamic environments than in static ones. Dynamic environments offer DMs the opportunity to correct poor decisions in response to unfavorable outcome feedback. Outcome feedback could also facilitate learning which improves the quality of decisions over time. These two considerations imply that DMs might perform better in dynamic situations than in static ones (Hogarth 1981).

In contrast, dynamic environments are often extremely complex, a feature that could diminish performance (Sterman 1989a, 1989b, Forrester 1971). Dynamic systems are characterized by networks of variables that form closed feedback loops. Often, the linkages between variables are non-linear and involve long lags between cause and effect. Several authors have suggested that the presence of feedback loops, time delays, and non-linearity can seriously limit DM learning (Sterman 1989a,1989b, Dörner 1980, Einhhorn 1980).

Anecdotal evidence from real-life business case studies also suggest that DMs can perform poorly in certain types of dynamic environments. For example, below, evidence will be presented that many firms seem to mismanage rapidly growing product markets. While there are many reasons why firms could fail in such situations, one possible explanation is that DMs do not make effective use of outcome feedback in growth situations.

Recently, a number of experiments in behavioral decision theory have dealt with dynamic environments (Kleinmuntz & Thomas 1987, Sterman 1989a, 1989b, McKinnon & Wearing 1985, Dorner 1980, Brehmer 1987, Diehl & Sterman 1994). These experiments have used a wide variety of methods, tasks, and subject pools and it is not possible to generalize about their results. However, many of the experiments have

shown that DMs perform very poorly in complex dynamic environments, especially when delays and feedback relationships are present (Sterman 1989a, 1989b, Dorner 1980, Brehmer 1987). In a review of studies of dynamic decision making, Brehmer finds that,

> "The most striking finding to emerge from the studies reviewed above is that people have considerable problems with time dependent processes. The subjects do not seem to use the historical information available to form the notion that the present state of the system is part of a lawful change. Instead, they use only the information at hand in deciding what to do. That is, the subjects do not seem to develop predictive mental models of these tasks, they only react by means of general heuristics. (Brehmer 1987)

Evidence from experimental economics suggests a more optimistic picture of DM rationality than do the results from behavioral decision theory. Experimental economics is concerned with the behavior of economic agents in market settings. Economic theory makes strong predictions about the market prices and quantities that result from the individual decisions of rational economic actors. A number of studies have demonstrated that economic theory does a good job of explaining the prices and quantities actually observed in laboratory markets. The accurate prediction of observed market behavior can be taken as indirect evidence that DMs are behaving rationally. As Plott states,

> "Even economists sometimes engage in disparaging remarks about the discipline because of doubts about the testability or the validity of the optimization hypothesis. If one looks at experimental markets for evidence, the pessimism is not justified. Market models based on rational choice principles do a pretty good job of capturing the essence of very complicated phenomena." (Plott 1986)

However, the experimental economics literature contains many results that cannot be explained by existing economic theory (Plott 1986). For example, a series of experiments has shown that asset prices often exhibit "price bubbles " that are not predicted by economic theory (Smith 1988)

As discussed above, psychologists and economists have begun to conduct experiments in dynamic environments. Generally, these experimental studies of decision making suffer from several limitations. These limitations are common to many of the experiments although not all of the problems apply to each study. The details of the individual studies will be discussed in chapter two.

1. The causal structure of the task environment is very simple. Relatively few experiments include the delays and closed feedback processes that characterize real systems.

2. The experimental situations are often very abstract. There is no way to relate the experimental results to real business situations faced by managers.

3. Feedback complexity, time delays and non-linearity are not varied as experimental treatments. This makes it difficult to attribute performance changes to the structure of the experimental environment

4. The effect of learning is not explicitly measured. Several of the experiments involve only single trials and do not allow for the possibility that, with experience, subjects learn improved decision rules.

This thesis presents an experiment that deals directly with the four limitations described above. In the experiments, subjects manage a simulated firm's investment, pricing and marketing decisions over the course of a product's life cycle. As will be discussed below, life cycle management is an extremely difficult managerial task and there are many real-life case studies where failure has resulted in huge financial losses.

Building on previous work, one possible explanation for these failures is that DMs mental models cope poorly with the feedback processes, time delays, and non-linearity that form the structure of the growth environment.

The market simulation used in the experiment is based on a well-established and realistic model of the product life-cycle that includes many of the feedback structures found in real markets. By varying parameters in the model, it is possible to evaluate how DM performance changes with the strength of important feedback loops or with the length of time delays.

The experimental results will be used to analyze two questions. First, how well do DMs perform in the dynamic product life cycle environment? To preview the results, the experiments show that DM performance is far from optimal even after DMs have had extensive experience with the task. In fact, after 200 quarters of simulated experience most DMs are outperformed by a simple decision rule.

Second, what accounts for the DMs poor performance? Following Sterman (Sterman 1989a, 1989b), it is hypothesized that the misperception of feedback relationships and time delays are responsible for the poor results. The experimental results support the proposition that the DMs learn to perform fairly well when important feedback links and time delays are unimportant. The same subjects perform poorly when feedback loops are strengthened and time delays are lengthened. More unexpectedly, the negative impact of feedback loop intensity does not diminish with experience.

The remainder of the thesis is organized as follows. Chapter two presents a review of the literature in behavioral decision theory and experimental economics that deals with dynamic environments. Chapter three deals with the problem of managing rapidly growing markets and discusses several case studies. Chapter four presents a model of the product life cycle and the structure of the simulation game. Chapter five explains the

experimental design, software and protocol and presents the results of the experiment. Finally, chapter six analyzes the results and presents conclusions.

Chapter Two
Review of Literature

This chapter reviews the experimental literature dealing with dynamic decision making. The experiments are drawn from the literature in behavioral decision theory, bootstrapping studies in management science, and experimental economics.

I. Behavioral Decision Theory

Experimental research in behavioral decision theory has identified numerous heuristics or rules of thumb that individuals use to simplify problems of judgment and choice (Hogarth & Makridakis 1981, Einhorn & Hogarth 1981, Slovic, Fischoff & Lichtenstein 1977). Psychologists argue that the use of these heuristics is pervasive because even relatively simple problems can overwhelm DMs' limited information processing capacity (Simon 1982). Faced with the necessity of making a choice, the DM has no alternative but to use a simplification strategy. Considerable evidence suggests that these heuristics are often dysfunctional and can produce systematic errors in judgment (Hogarth & Makridakis 1981, Einhorn & Hogarth 1981, Slovic, Fischoff & Lichtenstein 1977).

Behavioral decision theorists have begun to question the relevance of these findings to dynamic decision problems (Hogarth 1981, MacKinnon & Wearing 1985). In dynamic settings, there is two-way feedback between the environment and the decision maker (Sterman 1989a, 1989b). A DMs actions are affected by information about the environment. In turn, these choices feed back to change the state of the environment. This two-way interaction between the DM and the environment is an intrinsic part of many real-life situations.

Most of the experiments in behavioral decision theory require the subjects to make a single choice in a static environment. In these situations, no feedback is present. The

subjects' choices have no impact on the problem environment and, further, the DMs have no opportunity to revise their choices based on previous outcomes.

The static settings used in experiments are very different from many real-life situations in which DMs make choices and receive feedback about the effects of their actions. It is unclear whether the pessimistic view of human judgmental ability implied by static laboratory experiments can be generalized to dynamic environments. As Hogarth states,

> "Although recent research has identified systematic dysfunctional consequences of judgmental heuristics, the conditions under which heuristics can be valid have not been specified. In particular, whereas judgement is part of a continuous interactive process that people use to cope with the environment, most judgment research has focused on discrete incidents. This has led to underestimating the importance of feedback." (Hogarth 1981)

Surprisingly, very little research has dealt with decision making in dynamic environments. In a recent review Brehmer finds that,

> "There is surprisingly little research on the problem of how people learn to control complex systems. Studies in cognitive psychology are of little help here because they tend to be concerned with very simple systems." (Brehmer 1987)

Although the amount of research is relatively small, there have been a number of studies in dynamic decision making. The early experiments attempted to explain decision making in terms of the optimization model (Rappaport 1975). The experimental systems usually were very simple, often only a single state variable, so that optimal solutions could be derived. Performance was measured by comparing subjects' performance to optimal performance and deviations from the optimal solution were reconciled by adding constraints to the optimal model. One common constraint was to assume that the DM only looked one period into the future (Rappaport 1975). The solution to the constrained optimization problem was then compared to the subjects' decisions.

The difficulty of deriving optimal solutions restricted the early experiments to very simple environments (Hogarth and Makridakis 1981). Often they consisted of a single feedback loop and generally did not include non-linearity or time delays. Consequently, the subjects often learned to perform the tasks at close to the optimal level (Ebert 1972). As Rappaport states, "The experiments by Ray, Rappaport, and Ebert show that a substantial proportion of the subjects performed very efficiently, and that efficiency improved rapidly with practice." (Rappaport 1975)

More recently, a number of studies have used computer simulation to create more complex experimental settings that include feedback loops, time delays and non-linearity (Funke 1991, Brehemer 1992). A great deal could be said about each of these studies, but this review will focus on two specific issues. The first question is, how well DMs perform in dynamic environments. In analyzing this question, both first trial performance and learning over subsequent trials will be considered.

The second issue is what accounts for either good or bad performance. Several studies have found that DM misperception of time delays and feedback processes has been responsible for poor performance. Consequently, this review will focus on evidence about how these elements of the environment affect DM performance.

Simulation Studies in the Federal Republic of Germany

A number of researchers in the Federal Republic of Germany (FRG) have pioneered the use of computer simulations to study problem solving in complex dynamic environments (Funke 1988, 1991). These studies have considered a variety of dynamic problems ranging from relatively simple inventory control tasks to the management of the world ecology.

Some general themes concerning the causes of poor performance in complex dynamic settings have started to emerge from this research. Several of the studies found that task performance degraded significantly when there were significant time delays in the system, the variables were tightly interconnected, the relationships between variables were non-linear, and important system variables exhibited exponential growth. Two studies which illustrate these findings, Kluwe 1986 and Dörner 1980, are discussed in more detail below.

Kluwe 1986

The experimental setting was a system of 15 abstract variables that were not specific to any real-life setting. The subjects' task was to keep the variables close to fixed targets by issuing commands that added or subtracted from the values of the variables. The subjects were able to see only eight of the fifteen variables.

In one version of the task, there were no interconnections between the variables. A command affected a particular variable and there were no side effects on other variables. In this setting, the subjects learned to control the task effectively after 200 trials.

In other versions of the task, the variables were interconnected to form "causal nets". Kluwe reported that learning was much slower when variables were interlinked. It appears from the reported results that the causal nets did not form feedback loops.

Dörner 1980

The experimental setting was a computer simulation of a small town named "Lohhausene". The subjects acted as the mayor of the town for ten years and managed variables such as the production of the town factory, tax rates, and the number of

teaching positions created. The model consists of over 2000 variables and includes numerous feedback processes and delays.

Dörner found that the subjects' performance had wide variance: some people managed the town very well while others did very poorly. According to Dörner, three factors were responsible for poor performance. First, the subjects focused on the current state of the system and did not understand the processes that were generating important trends. In other words, they were unable to develop effective mental models of the system.

Second, the subjects could not deal with the rapid changes caused by exponential growth. Since the structural cause of exponential growth is positive feedback, Dörner's finding is consistent with previous studies that have implicated the "misperception of feedback" as a basic cause of poor performance. It is also consistent with previous studies that have show that subjects systematically underestimate the speed of exponential growth (Wagenaar & Sagaria 1975).

Third, the subjects thought in "causal series" rather than in "causal nets". Thinking in "causal series" means that people focused on the direct impact of their choices. Thinking in "causal nets" means that individuals considered the side effects of their decisions on other important variables. Dörner found that the failure to consider these side effects was an important cause of poor performance in Lohhausene.

Sterman 1989a

The experimental setting was a simulated economy in which the subjects made capital investment decisions. The simulation included the multiplier-accelerator mechanisms that are commonly used in economic models of investment behavior. The subjects played the role of the manager of the capital sector of the economy and controlled the system by ordering production capacity.

The simulation model is small, consisting of only three state variables, but includes several important feedback loops and time delays. For example, there is a delay of variable length between orders and shipments of capital goods. There is also a self-reinforcing feedback loop in which additional orders for capital cause the capital producing sector to expand by placing still more orders for their own output. Formal analysis of the simulation model has shown that this "self-ordering" process is extremely destabilizing when combined with a delay in the acquisition of capital (Sterman 1988b).

Sterman found that the average subject fell far short of optimal performance. In terms of satisfying demand for the product, the average subject generated costs that were thirty-one times greater than the optimal rule.

Sterman attributes the subjects' poor performance to two factors. The subjects' first problem is a serious misunderstanding of the delay between placing and receiving orders for capital. In the model, orders for capital that have not been delivered accumulate in a supply line. Sterman statistically estimated the decision rules that the subjects used to determine orders for capital. The magnitude of the supply line was one of the variables included in the estimated rule. The coefficient on the supply line usually was very small and often zero. Apparently, most subjects' decision rules ignored orders that had been placed but were still in the supply line. This caused gross over-ordering and enormous excess capacity. Subjects who took the supply line into account (larger coefficient on the supply line) did a much better job of balancing supply and demand.

Second, the subjects misperceived the effect of the positive feedback loop that drives self-ordering. When the self-ordering loop was eliminated from the model, the subjects' estimated decision rules performed at near optimum levels. This result suggests that the decision rules are appropriate for an environment without the self-ordering feedback loop but become dysfunctional when the loop is activated.

In a further analysis of Sterman's data, Bakken (Bakken 1989) presented some evidence that subjects' verbal understanding of the game is correlated with performance. He found that high performing subjects were more likely to mention the concepts of positive feedback and the system's steady state.

There are two important limitations to Sterman's experiment. First, the subjects played only one trial of the game. It is possible that subjects could learn to account for delays and feedback loops and improve their performance over time. Other studies of dynamic decision making have found significant learning effects (Brehmer1992).

Second, the effects of feedback loops and time delays on performance are analyzed indirectly through the subjects' estimated decision rules. This is a reasonable procedure but its validity depends on the validity of the specification and the functional form of the rule. Although the statistical measures (R2, standard errors) are reasonable, it is not known whether a small change in the model's specification would have given substantially different results (Leamer 1983). A more direct test would be to vary the strength of the time delays and feedback loops experimentally and directly observe the change in performance.

Sterman 1989b

The experimental setting was a simulated production-distribution system, the "Beer Distribution Game", in which the subjects played the role of inventory managers. Each subject was part of a team that included a retailer, wholesaler, distributor and factory. The subjects' task was to minimize inventory and stock-out costs by placing orders with their upstream supplier. For example, the retailer ordered product from the wholesaler and the wholesaler ordered from the distributor. The retailer, who is the first step in the supply chain, received a predetermined stream of customer orders.

The experimental system is very complex consisting of 23 state variables and numerous feedback loops, delays, and non-linearity. The key aspect of the system's structure is the variable delay between ordering and receiving product. For example, the retailer's delay in acquiring product depends on the ordering decisions made by all the other suppliers in the system. In the game, this delay ranges from four weeks to several months.

Just as in the previous study, the subjects performed very poorly. The average team's cost was ten times the benchmark cost with individual players exceeding benchmark costs by a factor of seven to thirteen.

Poor performance is attributed to two factors. First, just as in the previous study, the subjects mismanage the supply line by inexplicably ignoring unfilled orders. This leads to excess inventory and extremely high holding costs.

Second, the subjects tended to blame instability on external factors. After they played the game, the subjects were asked to sketch the pattern of customer orders that they thought was received by the retailer. In reality, there is a one-time increase in orders from four to eight. Most players drew fluctuating patterns and apparently believed that instability in the game was driven by an exogenous customer order stream.

Sterman argues that the subjects failed to learn effective decision rules because the focused on "open loop" explanations for the system's behavior. As Sterman states,

> "Most subjects attribute the cause of the dynamics they experienced to external events. Most blame their own poor performance on what they see as a perverse pattern of customer demand: the customers increased their demand, encouraging them to order additional beer, but suddenly stopped ordering just when the tap began to flow.... Most subjects attribute the dynamics to external variables which they believe to be closely correlated in time and space with the phenomenon to be explained. These explanations reflect an "open loop" conception of the origin of the dynamics, as opposed to a mode of explanation in which change is seen arising from the endogenous interactions of decision makers with their environment. Learning from experience may be hindered by such misperceptions of the origins of dynamic behavior. The erroneous open-loop

> attribution of dynamics to exogenous events thus draws subjects' efforts to learn away from the high leverage point in the system (the stock management policy) and towards efforts to anticipate and react to internal shocks." (Sterman 1989b)

The same limitations of the previous study also apply to this experiment. First, learning is not measured and it is possible the subjects would have learned to control the system effectively with additional practice. Second, the effect of delays is analyzed indirectly through the subjects' estimated decision rules. Again, while the estimated decision rules are generally very reasonable, it is not known how robust the rules would be to a specification change.

Diehl & Sterman 1994

The experimental setting was a simple inventory management system in which the subjects made production decisions over time. The subjects' objective was to minimize cumulative cost where cost depended on inventory availability and changes in the production rate.

Two treatment variables were manipulated to test the effect of delay and feedback loop intensity on the subjects' performance. The first treatement variable was the length of the production delay. It was hypothesized that a longer production delay would increase the subjects' cost relative to both an optimal and a naive "no control" benchmark. The second treatment variable was the strength of a feedback loop that connected shipments to demand. It was hypothesized that increasing the strength of the feedback loop would increase cost relative to both benchmarks.

Several interesting results were reported from this study. First, on the average, the absolute level of the subjects' costs were much higher than the optimal benchmark. Second, in the difficult delay and feedback conditions, the subjects' costs were higher

relative to the naive benchmark. In fact, in the most difficult scenarios, the "no control" benchmark outperformed the average subject. Third, the gap between the optimal decision rule and the rules the subjects actually used increased in the more difficult scenarios. The negative impact of delay and feedback loop intensity on performance is consistent with Sterman's findings discussed above. Finally, the subjects' performance relative to the optimal strategy did not vary with the difficulty of the scenario. This result was attributed to the fact that, for this particular task, deviations from the optimal strategy do not cause large increases in cost.

Brehmer 1990

The experimental setting was a computer simulation of a fire fighting system. The task was to dispatch and control fire fighting units in order to minimize the area burned. The subjects controlled as many as eight fire fighting units to put out two separate fires.

The computer simulation is very complex and includes numerous feedback loops and lags. For example, there is an important self-reinforcing loop in which an increase in the area burning causes the fire to spread even faster. There are also delays in moving the fire fighting units and in acquiring information about their position and effectiveness.

The presence of two distinct types of delays was a treatment variable in the experiments. The first type of delay is in moving the fire fighting units from place to place. Brehmer finds that the subjects learned to compensate for the delay by responding "rapidly and massively" to a fire (Brehmer 1990). On the first trial, the subjects' performance were about 8.5 times worse than optimal. Over six trials, performance improved by about 63%.

The second type of delay was in reporting the positions and actions of the fire fighting units. In this situation, it was much more difficult to learn to control the system. On the

first trial, the subjects' performance was nine times worse than optimal. Over six trials, performance improved by only 6%.

Broadbent and Aston 1978

The experimental setting was an econometric model of the British economy. The subjects, working in teams of 11, made periodic decisions for three control variables, government spending, tax rate, and the stock of money, in order to maximize a welfare function.

The system is fairly complex. The model consists of 22 equations and probably includes feedback loops although they are not explicitly reported. There are also explicit time delays between the variables. For example, there is a delay between an increase in government spending and higher inflation.

In one task, the subjects were asked to predict the impact of their decisions on inflation and unemployment. Over time, the subjects learned to predict unemployment accurately. In contrast, the subjects' predictions of inflation did not improve with practice. Broadbent and Aston argue that subjects failed predict inflation accurately because there were lags between changes in the control variables and their effect on inflation. However, these lags and their effects on learning were not varied as experimental treatments.

MacKinnon and Wearing 1980

The experimental setting was a computer simulation of a welfare administration project. The subjects allocated resources between different regions and methods of assistance. The objective was to maximize the average total welfare of the regions.

MacKinnon and Wearing experimentally varied the complexity of the system. Complexity was defined in terms of three dimensions: the number of elements in the model, the number of interconnections between elements, and the presence of random variation. Each dimension of complexity had two levels (eight total cells) and four subjects were assigned to each cell.

The dimensions of complexity were varied as follows. The number of elements was increased by having three methods of welfare assistance available instead of one. The number of interactions was increased by allowing migration between the three regions. Finally, randomness was introduced at meaningful points of the model.

On the average, the subjects performed at about the same level as a very simple benchmark. The benchmark was an equal allocation of resources between regions and methods of assistance. This implies that the returns to human decision effort were low.

Somewhat surprisingly, Mackinnon and Wearing found that increasing complexity did not reduce performance. Although it is unclear why performance was not degraded by complexity, they speculated that,

> "The complex situation may be different from the simple one. Heuristic methods which may be applicable to and successful in dynamic systems may not be applicable to the static decision-making tasks that have been widely studied. The utility of heuristic methods is suggested by their usefulness in solving complex dynamic problems such as playing chess." (MacKinnon & Wearing 1980)

One possible explanation for the MacKinnon and Wearing result is that added complexity does not necessarily degrade performance unless feedback loops and time delays also are added. While it is not possible to tell from the reported description of the task, the additional structure added to create the more complex case did not appear to add any important feedback processes.

Mackinnon and Wearing 1985

In this study, the task was to control a simple abstract system that consisted of one state variable and one negative feedback loop. The objective was to hold the state variable close to its goal by controlling a single input variable.

Mackinnon and Wearing found that the subjects outperformed a simple baseline strategy and that their performance improved over successive trials. The study concluded that,

> "The basic homeostatic mechanism is not likely to cause subjects difficulty in dynamic decision making tasks. Failure in more complex tasks can thus be attributed to the complexifying elements of these tasks." However, the dynamic complexity of the task is very low. The task is a first order linear system and is equivalent to filling a bathtub to a specified level with the drain open.

Hogarth and Makridakis 1981

The experimental setting was an extremely complex marketing simulation game called "Markstrat". The objective of the game was to maximize market share and cumulative marketing contribution by manipulating variables such as advertising spending, R&D, production, and market research. The subjects played the game in teams of seven.

The goal of the experiment was to compare the performance of the human teams with a set of arbitrary decision rules. For example, the simulation game required each team to determine the marketing budget. The arbitrary decision rule set the marketing budget at 10% of last period's revenue.

The results showed that the arbitrary rules performed as well as many of the human teams. As Hogarth and Makidakis state,

> "Probably the simplest way of summarizing the overall results is to state that the experimentally managed teams operated at approximately the same level as half the human teams. The surprising fact is that they did not perform significantly worse than their human counterparts." (Hogarth and Makidakis 1981)

Kleinmuntz and Thomas 1987

The experimental setting was a computer simulation of a medical treatment system. The subjects play the role of a doctor who must select tests and treatments for a series of patients. The setting is fully dynamic. The patients' health is assumed to decline over time so a prompt cure is essential. The wrong treatment or waiting too long to apply a treatment can kill the patient.

The risk that a given treatment will kill the patient was varied experimentally. Through computer simulation, it was determined that in the high risk situation a decision strategy that uses a significant amount of testing is best. In the low risk situation, a strategy of selecting treatments quickly without ordering any tests performs at close to the optimal level.

One interesting result is that in the low risk situation, only 8% of the subjects outperform a decision rule that randomly selects treatments without any testing. Kleinmuntz and Thomas observe that individuals may lack the cognitive resources to select the correct strategy in a complex dynamic environment.

In addition, some learning was observed. Over 30 trials, the subjects used fewer tests and treatments and took less time to cure a patient while the proportion of patients cured increased very slightly.

Conclusions

The first question addressed in this review is how well DMs performed in dynamic environments. In most of the experiments, first trial performance is far from optimal (Sterman 1989a,1989b, Brehmer 1989) and, in several experiments, the subjects did worse than simple arbitrary decision rules (Mackinnon & Wearing 1980, Sterman 1989a, 1989b, Hogarth and Makridakis 1981, Kleinmutz & Thomas 1987). Apparently, the subjects' previous experience did not equip them to control these systems effectively.

The evidence concerning learning is mixed. In several of the experiments, performance improved significantly over several trials (Brehmer 1989, Kluwe 1986, Broadbent & Aston 1978). In other experiments, the learning effect was very small (Brehmer 1989, Kleinmutz and Thomas 1987).

There are many possible reasons for the subjects' poor performance including confusing information displays and unfamiliarity with the experimental environment or the task. Several of the studies suggested a more systematic process in which poor performance was caused by the interaction between the DM and the feedback structure of the environment. Specifically, it was found that the misperception of time delays and feedback loops were an important cause of poor performance and also reduced the amount of learning (Sterman 1989a,1989b, Sterman & Diehl 1994, Brehmer 1989, Kluwe 1986, Dörner 1980, Broadbent & Aston 1978).

The still fuzzy picture emerging from the dynamic decision making literature is a continuum in which individual performance becomes progressively worse as feedback and time delays are added to a system. Decision makers' mental models function relatively well when they interact with environments that consist of simple cause-effect relationships. The same mental models seem to break down when confronted with the added complexity of feedback loops and time delays.

II. Experimental Economics

The objective of experimental economics is to test whether economic theory makes accurate predictions of market prices and quantities in various institutional contexts (Plott 1986). Economic theory is built from the assumption of individual rationality. Individuals are assumed to make choices that maximize a given objective function. These maximizing individuals interact in markets that have various institutional features. For some types of markets, economic theory makes strong predictions about the resulting market prices and quantities.

The focus of experimental economics has generally been on market outcomes as opposed to individual rationality. Researchers in experimental economics acknowledge that research in behavioral decision theory has shown that individual behavior often violates the rationality assumptions of economic theory (Plott 1986). However, this result is not thought to be relevant given the objective of understanding market behavior. From the viewpoint of experimental economics, the interesting question is whether or not economic theory, which is based on optimizing assumptions, makes better predictions than alternative models. As Plott states,

> "If the only question posed is, rational choice true or false?, then the answer is clearly false. Many critics of economics have claimed that the discipline is built on untestable foundations. Economists are indebted to psychologists for debunking such critics. However, the gratitude can go only so far. During the process of demonstrating testability, the psychologists disconfirmed the theory. The rejection of a theory of markets on the terms described above is not an especially challenging research objective. Those who study experimental markets tend to pose the questions in different ways. Rather than inquire whether a theory is true or false, they ask if the magnitude of error in the predictions of market phenomena is acceptable; or if no concept of degree of acceptability is readily available, the question becomes which of several competing models is the most accurate, fully realizing that the best model might still be poor." (Plott 1986)

The important point is that the unit of analysis in experimental economics is the market. Researchers in experimental economics do not claim that success in predicting market outcomes is evidence of individual rationality. Rather, the relevant question is how individual irrationality could affect market outcomes (Plott 1986). As Plott states,

> "When confronted with data that suggest the existence of erratic or irrational individual behavior, the implications are immediately evaluated in terms of the possible implications for a market level of analysis. Of course, when unusual market behavior is observed, one might then turn to models of irrational behavior to see if they contain the seeds of an explanation." (Plott 1986)

Economists believe that market processes can eliminate most of the effects of individual irrationality (Camerer 1987). For example, if they have sufficient resources, a small number of rational agents could force market prices to the efficient level that is predicted by theory (Camerer 1987). Similarly, irrational agents could learn from the behavior of more rational agents (Camerer 1987).

Researchers in experimental economics have considered many different market structures. In the early work, experimenters considered the simplest competitive market structure in which the subjects' supply and demand schedules were fixed (Smith 1982). These experiments focused on the question of whether price and quantity converge to the static equilibrium values predicted by economic theory.

One of the most studied market institutions is the double auction. In a double auction, subjects are divided between buyers and sellers. Each buyer is provided with a marginal value for units of the product and each seller is provided with a marginal cost. The buyers enter bids and offers to buy and sell units of the commodity. Trades are executed whenever a bid made by a buyer is accepted by a seller.

Numerous replications using the double auction have shown that prices usually converge to the price predicted by economic theory (Smith 1986). This result has survived numerous variations on the basic market organization such as adding

middlemen who buy from producers and resell to buyers, or allowing demand to follow a cyclical pattern. Also, several experiments have confirmed the predictions of economic theory in more complex double auction markets in which two substitute commodities are traded. Speaking of these results, Smith finds that,

> "The most important feature of these results is that the high predictive power of demand theory in the single market setting extends to the interdependent two-market setting. Subject traders, in effect, solve a set of simultaneous nonlinear equations using the double auction trading rules, without being aware that this is the market consequence of their behavior." (Smith 1986)

A second market environment that has been studied extensively is the posted price environment (Smith 1986). In this setting, sellers post prices which they cannot change until the next time period. Buyers decide how much to purchase from each seller at the seller's posted price. After all the trades have been completed, the sellers post the next period's prices.

Extensive replications of the posted price institution have shown that market prices are close to those predicted by economic theory (Smith 1986). However, prices are higher than in the double auction markets, a result not predicted by economic theory. As Smith states,

> "Even now no theory about the influence of the posted price institution has been published to my knowledge, but the effect has persisted under a variety of parametric situations." (Smith 1986)

A third set of experiments have dealt with asset pricing (Smith 1986, 1988). In these experiments, subjects trade an asset in a double auction market over several time periods. The dividend is uncertain but has a distribution that is known to the subjects. These experimental markets are more complex than the double auction markets because they also allow for capital gains when asset prices change. The possibility for capital gains adds a dynamic element that complicates the subjects' decision problem. Traders must

anticipate how prices will change over time and their demand schedules will depend on that expectation.

Economic theory predicts that in a market populated by rational economic actors, the asset's price should equal the present value of its expected net cash flow. In these experiments, expected net cash flow equals the expected value of the dividend.

In many of these experiments, the asset's price converged to its predicted level only after a transient price bubble much like the speculative booms and busts observed in the real-world (Smith 1986,1988). In the long-run, after the bubble dissipates, price converges close to the value predicted by economic theory. The price bubble is not predicted by the theory and could be interpreted as a case where individual irrationality affected a market outcome.

The experimenters argue that price oubbles could be consistent with rational behavior if the subjects had different expectations of future prices. Even though subjects had common expectations of future cash flows, price expectations could vary because of uncertainty about how other subjects would behave.

Overall, the experimental evidence concerning market efficiency is mixed. In some settings, such as double auctions, prices generally converge to the level predicted by economic theory. In other settings, such as asset markets, irrational behavior was not eliminated by market processes (Smith 1988, Camerer 1987).

Overall, there is some evidence that the more complex markets generate worse performance than the simple markets. The simple double auction markets generally converge to the equilibrium predicted by economic theory. In the more complex asset markets, prices often deviate significantly from theoretical predictions (Smith 1988, Camerer 1987). As discussed above, a similar relationship between complexity and performance has been found in dynamic decision making studies conducted by psychologists.

The market studies from economics may have limited applicability to the experimental environment used in this thesis. In the discussion of the literature from behavioral decision theory, it was argued that time delays, feedback loops, and non-linearity caused poor performance. None of these factors are present in the market studies. The finding that subject's behavior is usually consistent with economic theory in simple markets does not imply that markets would produce rational behavior in more realistic settings (Smith 1986).

Kampmann examined the impact of alternative price setting regimes in a market that included time delays and feedback processes (Kampmann 1992). Kampmann found that the presence of feedback loops and time delays lowered performance, a finding that is consistent with previous studies in behavioral decision theory. Unlike previous studies, Kampmann's subjects competed in an experimental market, allowing him to test whether competition would eliminate the negative effects of dynamic complexity. The results demonstrated that price competition eliminated only a small percentage of the cost of dynamic complexity.

III. Bootstrapping

There is a third body of literature that has dealt with decision making in dynamic environments. Early in the evolution of decision making research, Bowman classified errors in human decision making into two types, variance and bias (Bowman 1963). Variance is the inconsistent use of a decision rule that on the average is reasonable. Variance is caused by random errors in human judgement.

Bias is the error caused by "Misperceiving the relative importance of the system variables in making decisions" (Bowman 1963). A biased decision rule applies incorrect weights to the available information cues.

Bowman and subsequent researchers argued that variance is by far the most costly source of error. This implies that, at least in the context of the study environments, DMs use effective rules. The problem is that they do not apply the rules consistently. A set of techniques known collectively as "bootstrapping" have been developed to estimate statistically the DMs rule (Armstrong 1985). Once the rule is known, it can be applied mechanically to make judgments or forecasts (Armstrong 1985).

However, most of the bootstrapping experiments have used relatively simple settings. Many of the experiments have involved discrete tasks in static environments. Relatively few experiments have been conducted in the type of dynamic situation discussed by Hogarth. As Makridakis states,

> "In summary, the literature indicates that inconsistency reduces the quality of decisions in the short-term when environmental conditions are stable and non-competitive. However, little can be said about the more usual ecology of human decision making where sequences of decisions are taken over time in competitive and possibly unstable conditions. "(Makridakis 1981)

A few of the bootstrapping studies have dealt with dynamic environments (Ebert 1972, Moskowitz 1974, Bowman 1963, Remus 1973). Remus used a complicated business game in which subjects made pricing, marketing, and production decisions (Remus 1978). Remus found that both variance and bias were correlated with the subjects' performance relative to other players and the performance improved with experience in the game.

Moskowitz used a production scheduling task in which the subjects chose the levels of work force and production over time (Moskowitz 1974). Moskowitz found that performance was far from optimal and that inconsistency accounted for most of the difference.

In contrast, Sterman found that inconsistency accounted for some of the poor performance of the subjects in his studies. The bootstrapping improvement was only

about 5% in contrast to a better rule which could produce improvements of 1000% to 3100%. This finding may suggest that bias may become more important in more complex environments that include time delays and feedback.

IV. Conclusion

Compared to static environments, there has been relatively little experimental work done on decision making in dynamic settings and it is still premature to draw general conclusions. One finding that may be emerging from the literature is that DM performance depends on the dynamic structure of the environment. Specifically, in the relatively simple environments used in experimental economics, bootstrapping experiments, and the early work in dynamic decision making, DM performance was often good especially after experience with the task.

Several of the more recent experiments suggest that DMs have a great deal of difficulty managing dynamic systems that include feedback and time delays. The experiments discussed above indicate that humans don't perform particularly well in comparison to simple decision rules and that learning is often very difficult. However, because few experiments have focused specifically on learning, relatively little is known about a DMs ability to learn in complex dynamic environment.

Chapter Three
Case Studies in Growth Management

Most of the experiments in dynamic decision making have considered relatively simple static environments. As argued in chapter one, these settings generally do not include the feedback processes and time delays that characterize many real situations.

The settings used in the experiment reported here involve the management of a new product over its life-cycle. The simulated market includes numerous feedback loops and time delays that connect the market, a firm and its competitors. The simulation model is discussed in chapter four.

One advantage of using life-cycle management as the experimental setting is that this situation has a real-world analog. There are numerous examples of firms that have faced the problem of managing the transition from rapid growth to maturity. These examples allow an examination of how well firms perform the task in the real situation and, in some instances, the reasons for success or failure. The analysis of the real situation can be compared qualitatively to the subjects'performance in the experiment.

The remainder of this chapter is organized as follows. The first section reviews several background studies dealing with growth management. The second section discusses the past work in system dynamics that deals with growth management. The final section presents several case studies that deal with the particular issue of growth and saturation.

I. Background Studies

A number of studies have dealt with the issue of how well firms manage rapid growth. Birch considered a sample of several hundred high growth companies (Birch 1987) and found that,

> "We know that these fast growing companies are volatile. Growth looks much more like a roller coaster than a ramp. In our broader data-base, we find that the single best predictor of a company's growth up to a two to three year horizon is previous decline. The best predictor of decline is recent rapid growth." (Birch 1987)

In addition, Birch analyzed the survival profile of new businesses. He found that one-half survive their first five years and only thirty percent last into their fifteenth year (Birch 1987).

Clifford and Cavanaugh examined a sample of rapidly growing mid-sized companies (Clifford and Cavanaugh 1986). They found that

> "The midsized companies as a group were all over the lot. Of them, the great majority had earnings that grow very rapidly, decline severely or were negative. A similar analysis for 1978 to 1983 showed a comparable pattern. Even the winning performers we studied, the best of the middle-sized group are not immune from instability. At the extreme, about one-half of the chief executives told us that at some point in their history their companies had been "broke" or near "broke". Between 1964 and 1985, 85% of the companies experienced at least one decline in sales or earnings and most had more than one down year. More that a third showed a loss in at least one year." (Clifford and Cavanaugh 1986)

In an article entitled, "Emerging Growth Companies; Why So Many Peak So Early", Forbes magazine analyzed the performance that had previously made its list of promising growth companies. Forbes was led to ask, "Why does it happen? Why do so many small promising companies peak so early? Why doesn't success breed success?"

(Forbes 1985). The article provided nine examples of promising companies that grew very rapidly and subsequently collapsed.

Bayus studied sales, price, and cost data for 27 durable goods (Bayus 1993). In his study, he estimated a sales diffusion model and found that the products could be classified into five segments according to the pattern of sales growth and decline. Although the timing of the peak varied between the five segments, all five segments exhibited the overshoot and decline pattern.

Previous Work in System Dynamics

The studies presented in the previous section suggest that businesses in rapidly growing market suffer a great deal of instability. However, these studies do not establish the reasons for the instability. Certainly, there are many possible explanations for the growth and decline syndrome such as bad luck, organizational difficulties, and inadequate financial backing (Forbes 1985).

System dynamics studies have isolated several "generic structures" or "systems archetypes" that relate to the problem of growth management (Paich 1986, Senge 1990). These generic structures are analytic templates that help explain the problems experienced by firms across many different businesses.

One such structure is the growth and under-investment syndrome that has been prominent in many system dynamics studies (Sterman 1988, Senge 1990, Forrester 1968, Fortune 1985). One such study dealt with the case of the People Express Airline. People Express grew very rapidly for five years and then quickly collapsed into bankruptcy. Many explanations have been offered for the failure of People Express including the inefficiency of "soft" management practices and the development of yield management systems by the major airlines.

Sterman proposed a useful systems explanation for the collapse that emphasized the interaction between growth and under investment (Sterman 1988). The growth and under-investment analysis initially had been developed for a firm in a high- tech manufacturing business. It was discovered that the manufacturing firm artificially limited its own growth by failing to invest in adequate production capacity. Sterman found that the same basic process applied to People Express. Senge summarizes Sterman's argument this way,

> "People Express's chronic problems with service quality and having enough competent and committed service personnel suggest subtle similarities to WonderTech with its problems of inadequate manufacturing capacity and eroding delivery service - even though the specifics at People Express differed in almost every way from the specifics at WonderTech. WonderTech foundered because of worsening delivery times and eroding delivery time standards. People Express foundered because of declining customer service quality and standards for service. But despite all those differences, underlying both were the dynamics of growth and under investment, the systems archetype that explains one of the most common ways that organizations limit their own growth." (Senge 1990)

A second archetype also deals with the dynamics of rapidly growing markets. Markets often expand rapidly as word of mouth communication, investment in advertising, and falling prices attract new buyers. Eventually, the stock of new buyers is exhausted and demand consists mostly of replacement sales (Mahajan, Muller & Bass 1990). For durable goods, sales typically peak three to six years after the product is introduced (Urban & Hauser 1980).

During the transition to replacement demand, firms often suffer from excess capacity and low profits due to unrecovered fixed costs and falling prices. As Porter states,

> "As the industry adjusts to slower growth, the rate of capacity addition in the industry must slow down as well or over capacity will occur. Thus companies' orientations toward adding capacity and personnel must fundamentally shift and be disassociated from

> the euphoria of the past. Rapid growth will no longer quickly cover mistakes by rapidly eliminating excess capacity. These shifts in perspective rarely occur in maturing industries and overshooting of industry capacity relative to demand is common. Overshooting leads to a period of over capacity, accentuating the tendency during transition toward price warfare." (Porter 1980)

This period of industry turbulence is sometimes referred to as an "industry shakeout" and usually results in many firms leaving an industry (Wasson 1978, Klepper & Graddy 1990). For example, during the shakeout of the color television industry, thirteen of the nineteen firms failed to survive (Willard and Cooper 1985). A similar story can be told about video games (EIA 1984), hand calculators (EIA 1989), and home computers (EIA 1989).

These two system archetypes create a difficult managerial dilemma. On one side of the dilemma, an overly conservative policy of restricting capacity growth can lead into the trap of artificially limiting growth, market share and long-term profitability. On the other side, an overly aggressive policy can produce huge losses when the market inevitably saturates and declines. The difficult managerial task is to navigate a course between the two extremes.

II. Case Studies

Management of the growth-saturation dynamic will be the major focus of this thesis. Consequently, the next two sections describe a series of case studies dealing with this issue. The first group of case studies are meant to illustrate the operation of the growth-saturation process in several businesses. The last three case studies provide more detail about the strategies used by individual firms.

Tensor Lamps

Jay Monroe, the inventor of the high intensity designer lamp, saw sales of his product skyrocket. In fact, the market for high intensity lamps grew very rapidly for five years and then quickly collapsed. As shown in figure 1, Tensor sales reached $7,445,000 in 1966 and then fell to $3,262,000 in 1968.

Figure 1
Tensor Corporation
Sales & Net Income (Thousand $)

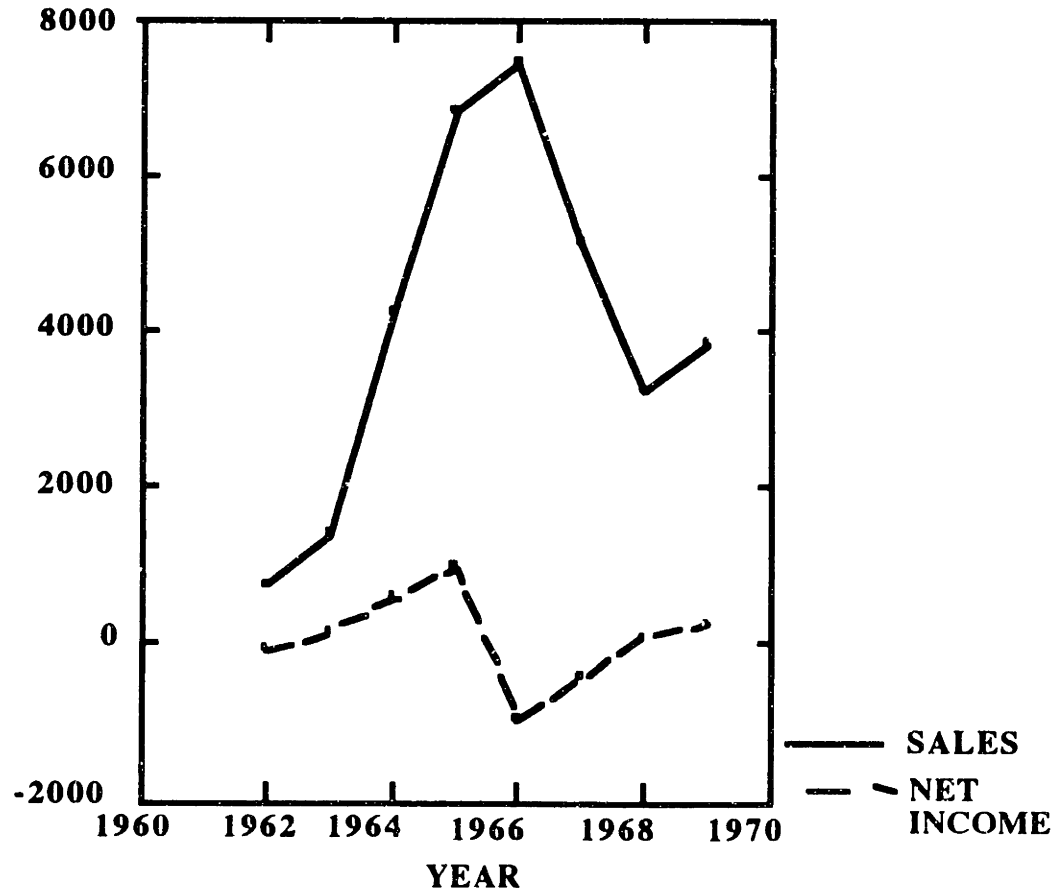

Tensor, the largest competitor with a 50% share, was nearly bankrupted in the process. Figure 1 also shows the pattern of Tensor's net income from 1962 through 1969. Tensor substantially expanded its production capacity in order to maintain market share and was left with enormous excess capacity when the market declined (Salter 1965).

Jay Monroe, the CEO of Tensor, attributed the collapse to one of the self-reinforcing processes described above. He believed that extensive advertising spending had pushed sales to unsustainable levels. Heavy advertising compressed demand into a few years which caused the large peak and subsequent collapse. As CEO Jay Monroe stated,

> "The second factor which affected our forecasts was the extra consumer demand generated by the large amount of advertising and promotion. We "over-advertised" to try to build up our name while the market was still immature, and this contributed to a larger final demand than would have been experienced in a time of less aggressive promotion ... These campaigns stimulated some demand which was not sustainable when the promotion level became more normal." (Salter 1965)

Snowmobiles

The snowmobile industry grew rapidly until 1971. As shown in figure 2, unit sales increases from 58,000 in 1968 to 342,000 in 1971 (ISIA 1984). Sales declined rapidly after 1971 and fell to 59,000 units in 1984 (ISIA 1984).

Figure 2
Snowmobile Sales
(Thousands of Units)

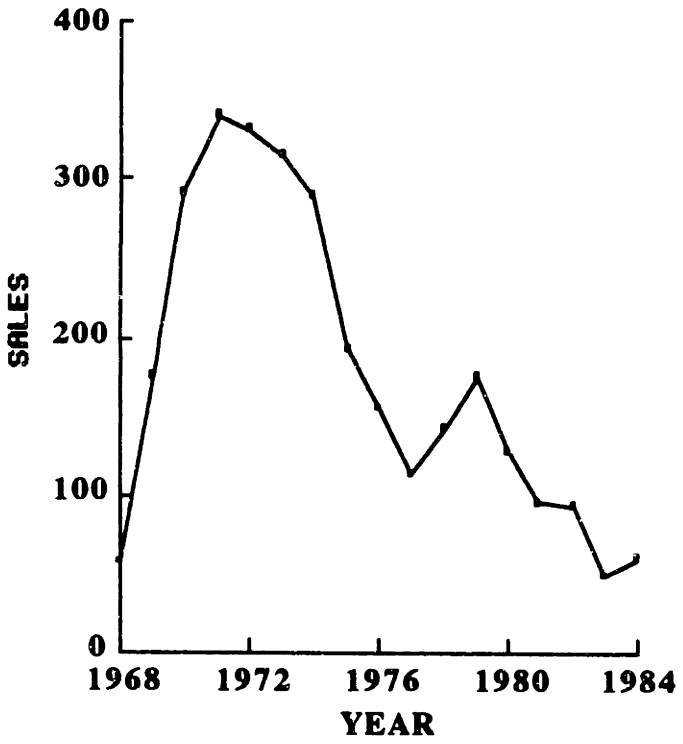

The firms in the industry were battered by the decline. The shakeout was so severe that ninety per cent of the firms in the industry before the decline were forced out of business (Business Week 1976). For example, Scorpion, one of the surviving firms, saw its net income fall from a peak of $1,234,000 in 1971 to a loss of $5,637,000 on $6,760,000 of sales in 1973 (Salter 1975).

Spokesman for the industry attributed the decline to a series of external events. A typical explanation was,

> "For North American snowmobile producers there have been few good moments since the winter of 1970-71 when customers bought 425,000 of their vehicles and clamored for more. Environmental

> objections, safety challenges, and then the recession dropped sales precipitously (Business Week 1976)."

Other explanations attributed the decline to the energy crisis and bad snow years (Salter 1975).

While they have some validity, all of the "external" explanations overlook some basic facts. First, demand peaked in 1971 which was before the energy crisis and the safety problems became important. Second, relative to the peak years from 1970-1973, demand remained low during the entire period through 1984. Economic growth was strong in some of these years yet there was no substantial rebound in sales. The same analysis applies to the "no snow" explanation.

Home Computers

The home computer market (computers priced at less than $2000) grew explosively beginning in 1982 (EIA 1989). Fueled by a price war between Texas Instruments and Commodore, unit sales increased from two million units in 1982 to over five million units in 1984 (EIA 1989).

Sales quickly declined between 1984 and 1986. As shown in figure 3, sales fell to 3.8 million units in 1986 (EIA 1989). Commodore, one of the industry's largest firms, suffered a fifty per cent sales decline from 1984 to 1985 (Wall Street Journal 1985).

Figure 3
Home Computer Sales
(Thousands of Units)

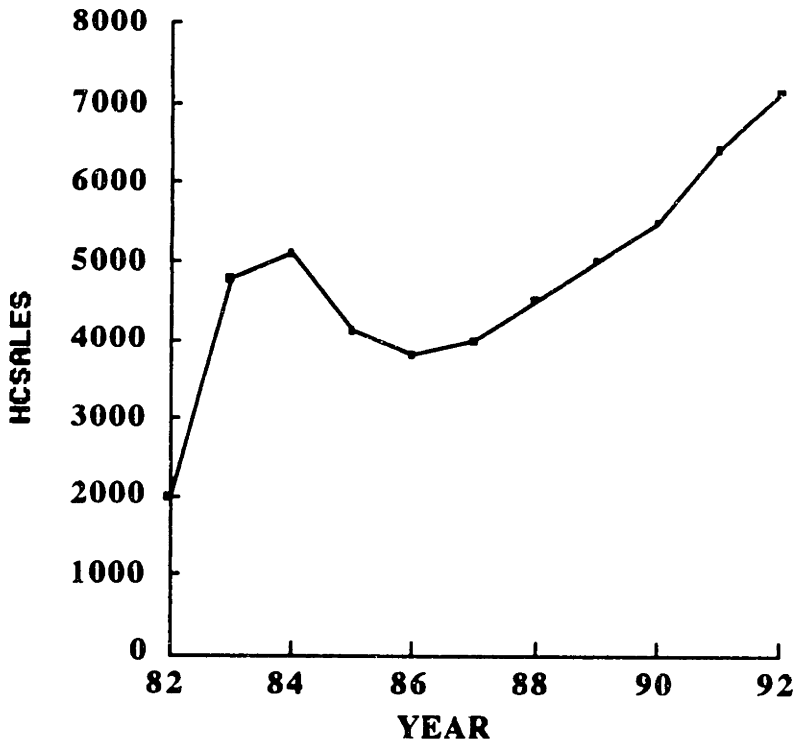

The sales decline devastated the home computer industry. Texas Instruments lost over $600 million in the downturn and was forced to exit the market (McClellan 1984). In 1985, Commodore was operating at less than 50% of capacity and lost more than $80 million during the fourth quarter of that year (Kneale 1985).

Toys and Games

Life cycle management is the critical problem in the toy and game industry. There are many examples of firms that rode a boom market for a single toy and then collapsed after

demand peaked including Coleco with Cabbage Patch Kids, and Worlds of Wonder (WOW) with Teddy Ruxpin and Laser Tag.

WOW's experience with Laser Tag and Teddy Ruxpin is a classic example of the boom-bust syndrome. WOW introduced laser tag with one million dollars of advertising; far more than toy makers usually spend (Business Week 1988). The advertising created enormous demand and Laser Tag became the nation's hottest selling toy. WOW had great difficulty increasing production and missed delivery dates were common. By the time WOW built adequate production capacity, demand had peaked and many retailers dumped Laser Tag at a loss. A similar process occurred with the Teddy Ruxpin talking bear (Beinhocker 1991).

The collapse of Laser Tag and Teddy Ruxpin bankrupted WOW which lost $187 million during the third quarter of 1987 and was forced to file for chapter 11. WOW left its creditors with $260 million of bad debts including $40 million owed to five Hong Kong manufacturers.

Tyco has taken a very different approach to life cycle management. Tyco extends the life cycle by intentionally restricting the supply of hot selling toys. For example, Tyco sold a popular doll named Oopsie-Daisy. Production of the toy was kept at low levels so that demand would last for several seasons (Popper 1989).

Instead of expanding capacity, Tyco has plowed its profits from popular toys into product development. They have built a diversified portfolio of toys in which no single product contributes more than 16% of total revenue. The strategy of limiting production and diversification has been very successful as earning and sales have both increased for five consecutive years (Roman 1989).

Consumer Electronics

In a variety of consumer electronics industries, firms have suffered from the growth and collapse syndrome. Recently, many Japanese companies significantly expanded capacity in several product lines only to see demand stagnate and decline (Schlessinger 1992). Two of the more well-known examples, VCR's and video games, are discussed in more detail below.

VCR's

Figure 4 shows the pattern of sales in the VCR industry. Sales grew rapidly until 1986 and then began the familiar pattern of decline. During 1989, sales fell about 10% from their peak (EIA 1993).

Figure 4
VCR Sales
(Thousands of Units)

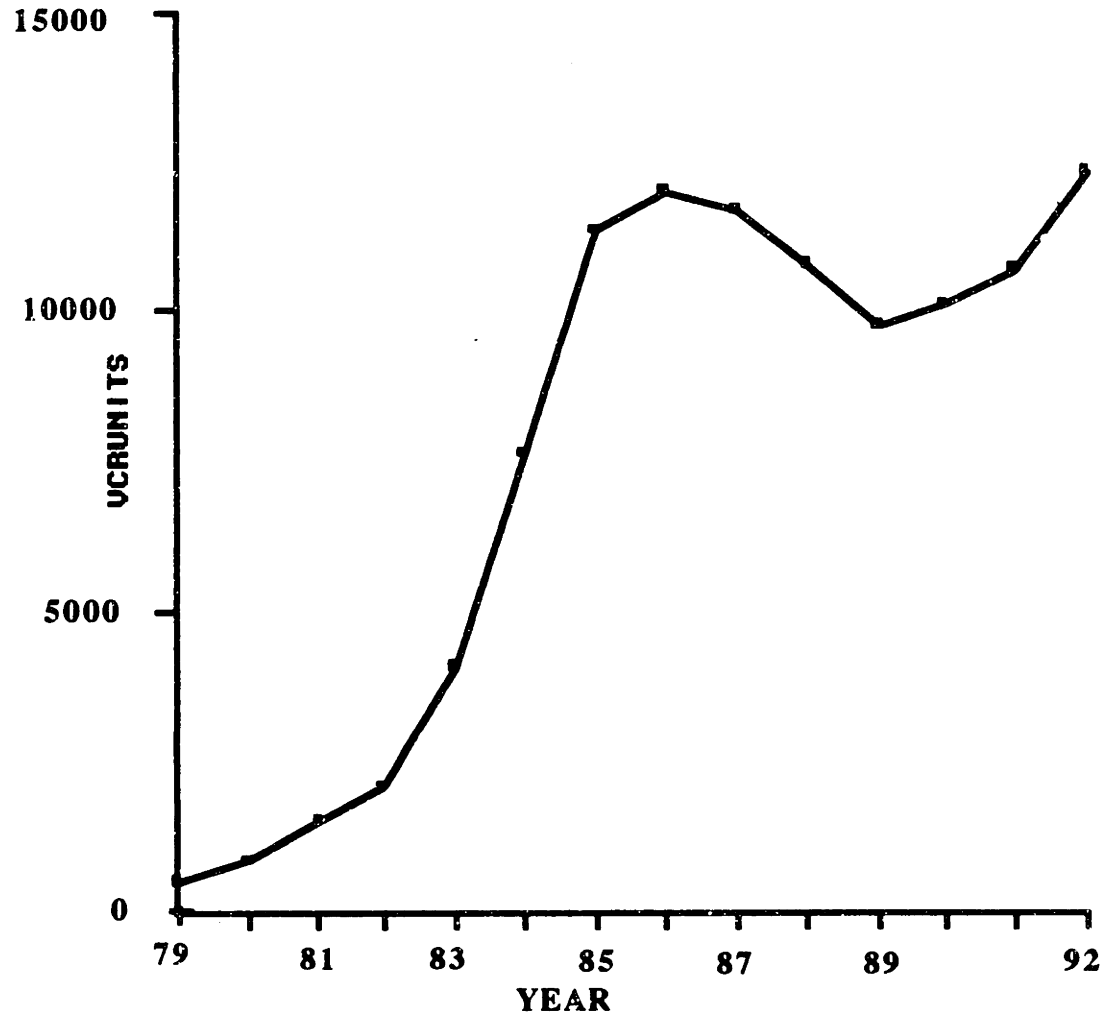

By 1992, the primarily Japanese producers were suffering from the typical symptoms of market growth and saturation. Prices were at half their 1980's peak and sales were down one-third from their high (Schlessinger 1992). As shown in figure 5, revenue corrected for inflation had declined precipitously (EIA 1993). As a result, the JVC corporation, which has been the leader in VCR production, was expected to report its first loss since 1951 (Schlessinger 1992).

Figure 5
VCR Real Revenue

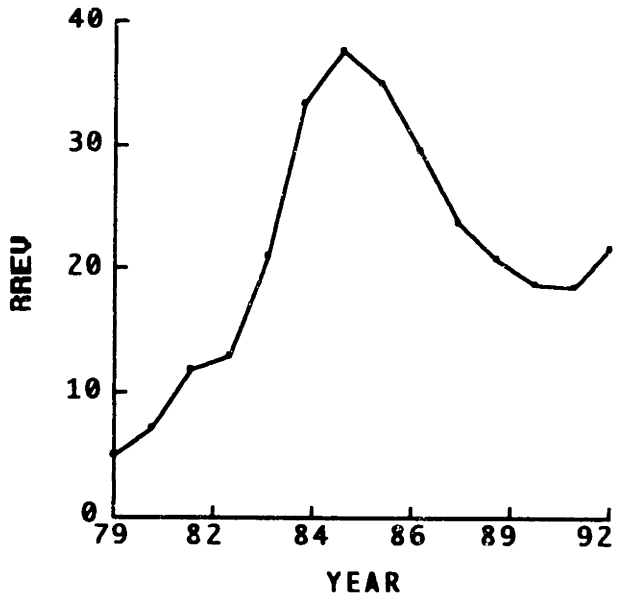

Video Games

The video game industry is a fascinating example of the market growth and saturation dynamic. Beginning in 1979, the demand for video game players and cartridges grew rapidly. Figure 6 shows the pattern of revenue beginning in 1979. In 1982-1983 total sales peaked at over $3 billion and then rapidly collapsed to $100 million in 1985 (Cowan 1990).

Figure 6
Video Game Revenue
(Billions of $)

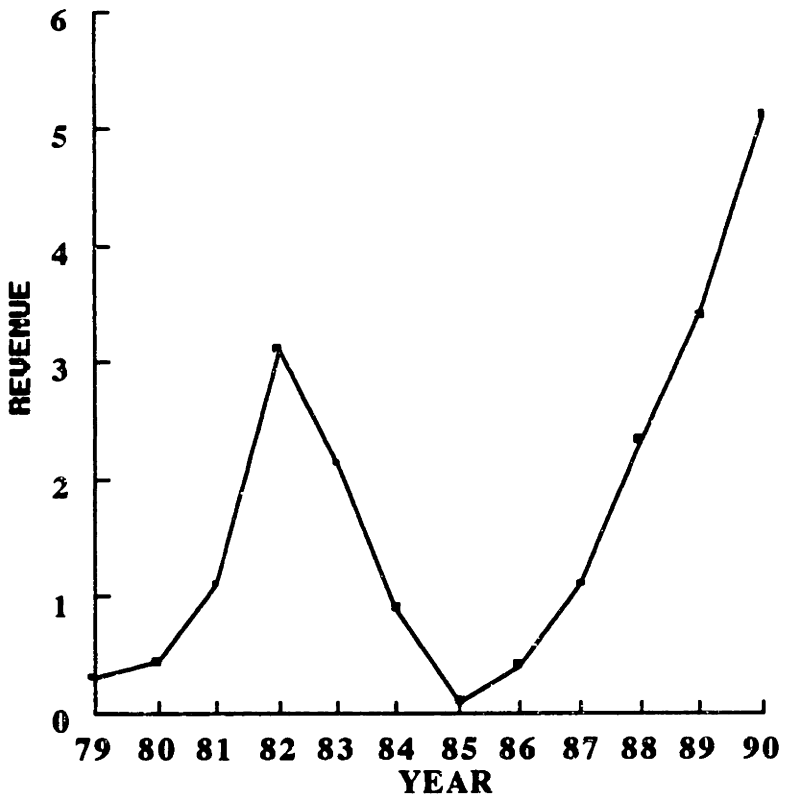

The growth and collapse of the market devastated the firms in the industry. Warner Communications bought Atari, a major player in the industry, in 1981 for approximately one billion dollars (Petre 1985). At the peak of the market, Atari earned an operating profit of over $200 million on revenues of $1.2 billion (Petre 1985). Following the market collapse, Atari lost $539 million in 1983 and was finally sold to Jack Trameil at concessionary terms in 1985 (Petre 1985). A smaller manufacturer of video games, Wells-Gardner, saw its net income rise from $410 thousand in 1979 to $6.5 million in 1981 (Angrist 1983). Wells-Gardner's sales fell 63% between 1982 and 1983, nearly bankrupting the firm.

Beginning in 1986, the video game industry rebounded. Led by the Japanese manufacturer Nintendo, combined hardware and software sales in the U.S. increased from $100 million in 1985 to $1 billion in 1987, to $2 billion in 1988 (EIS 1988), and over $5 billion in 1990. Some estimates put the global market value at over $10 billion in 1990 (Neff 1990). Sega, which beat Nintendo in introducing the next generation 16 bit game player, had captured 49% of the U.S. market and had grown from less than $1 billion in 1990 to nearly $4 billion in 1993 (Brandt 1994).

In early 1994, the video game market was again showing symptoms of overshoot and decline. The 16 bit generation of game players was starting to saturate causing inventories to build and prices to fall (Brandt 1994). The video game market as a whole peaked in 1990 and declined by over 10% between 1990 and 1991 (Pereira 1991). The major producers, Nintendo and Sega, have cut prices by up to 50% in order to stimulated demand (Pereira 1991). After growing at a compound rate of nearly 40%, Sega's sales were expected to decline slightly in 1994 and its profit to fall by 63% (Brandt 1994). Ninetndo's sales are also sluggish and its profit is projected to decline by 40%.

Video games could be viewed as just another fad market. Most video games are sold to children, who are notoriously fickle consumers. The cause of the market collapse could be viewed as an unpredictable change in children's tastes and attitudes. One typical explanation is, "It was Atari that launched the first video game boom in 1976 and saw it expand to a $3 billion business by 1982. Soon scores of companies rushed out games. Most were mediocre, though and jaded juveniles quickly lost interest." (Brandt 1990)

A more sophisticated explanation for the growth and collapse has important strategic implications. This explanation focuses on fundamental aspects of the market structure that systematically generate overshoot and collapse. Once these structures are understood, compensating strategies can be developed.

Word of mouth communication is an extremely powerful mechanism in children's products, as children are great imitators. In markets with a strong propensity for imitation, word of mouth communication creates a self- reinforcing process in which recent sales breed still more sales (Lilien 1992).

In addition, advertising creates a second self-reinforcing process. Video game manufacturers spend a great deal on advertising. During the early 1980's, advertising for personal computers and video games grew dramatically . More recently, Nintendo raised its advertising expenditure by 58% from 1989 to 1990 (Pereira 1991). Effective advertising leads to more sales which increases the cash flow available to fund still more advertising.

Finally, the presence of a learning curve creates a third self-reinforcing process. In many electronics businesses, costs fall substantially as cumulative production increases (Lilien 1992). Lower prices increase sales which further increases cumulative volume (Lilien 1992). In fact, the average price of a programmable video game fell from $116 in 1983 to $43 in 1985.

Market demand will peak and decline very quickly if these self-reinforcing process are strong. At any given time, there is a stock of potential buyers who are attracted to the product but have not yet purchased it. The stock of potential buyers will be exhausted quickly if the three self-reinforcing processes cause sales to grow rapidly. Depending on the strength of the loops, demand often peaks when market penetration reaches between forty and sixty per cent (Mahajan, Muller & Bass 1990).

The collapse in demand will be worse if the product is infrequently purchased because the stock of potential buyers will be replenished very slowly. In the case of video games, once someone bought a game player there was no reason to buy another until the next generation of players was introduced. Given the inherent delays associated with

research and development and retooling production facilities, there was ample time for the market to collapse between product generations.

Overall, two basic structural characteristics generated the behavior of the video game market. First, there were strong self-reinforcing processes that caused demand to grow rapidly and to exhaust the stock of potential buyers. Second, because there was no incentive to repurchase the product, the stock of potential buyers was replenished very slowly.

There are several strategies available for dealing with this situation. As discussed above, Nintendo led the revival of the video game industry and attempted to avoid the growth-saturation syndrome. First, Nintendo intentionally limited the supply of video game consoles and cartridges by restricting the production of Nintendo machines to Japan. (Hill 1989). Limiting supply should flatten the peak of demand and extend the life-cycle of the product.

Second, Nintendo raised prices to limit demand and relied on advertising and word of mouth to drive market growth (Hill 1989). This pricing strategy was complementary to the policy of limiting production because it rations supply by price instead of availability. Higher prices generate cash flow which can be used to fund research and development for new generations of games. The combination of the low production and high price enabled Nintendo to earn a pretax profit of $725 million on $2.7 billion of sales in 1989 (Neff 1990).

Finally, Nintendo consciously managed the growth-saturation dynamic. Nintendo developed a new and more powerful game console that uses a 16- bit microprocessor where the old console uses 8-bit technology. Nintendo believes that sales will begin to fall when market penetration reaches 40% and it is delayed the introduction of the 16- bit technology in the U.S. until market penetration of the old technology reached that level

(Neff 1990). Delaying introduction was designed to provide customers with an incentive to renter the market at the time when demand was beginning to fall.

However, despite its seemingly intelligent management of the growth-saturation dynamic, Nintendo's strategy was not as successful as it had hoped. Sega introduced a 16 bit game player before Nintendo and used its technological advantage to take a significant market share (Brandt 1994). Nintendo's policy of limiting supply angered retailers and made them receptive to alternative suppliers. A major toy distributor, Stephen Sandberg, stated,

> "When you try to control a product by not shipping enough of it, there's a thin line. You can absolutely kill a product by not shipping enough of it if retailers lose sales. Nintendo's tactic made retailers more apt to look to someone else". (Brandt 1994)

Chain Saws

The consumer chain saw industry during the 1970's and early 1980's is an excellent example of the growth management dilemma. The chain saw market is divided between the professional and casual user segments. The professional segment mostly consists of industrial users who purchase large, relatively expensive saws. The professional market is relatively stable and has grown steadily with the exception of recession years.

The casual user segment grew explosively during the 1970's and then collapsed in 1980-1981 (Porter 1981). As shown in figure 7, chain saw sales peaked at nearly three million units and then declined to 1.6 million units in 1985.

Figure 7
Chain Saw Sales
(Thousands of Units)

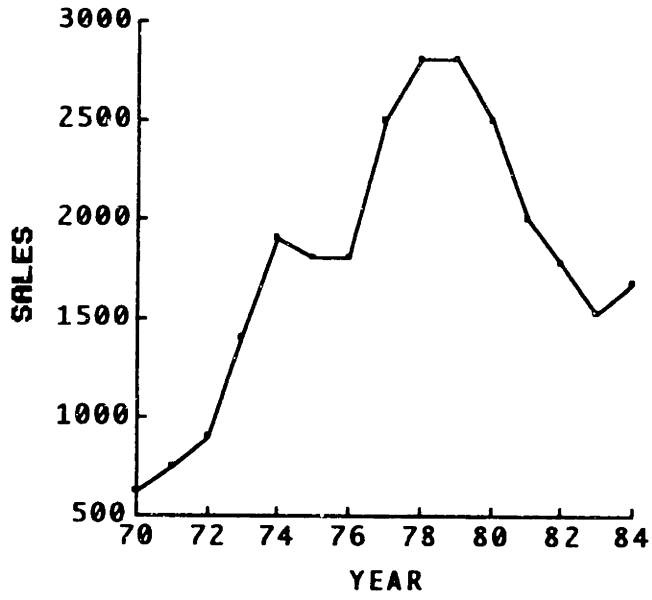

Market growth was driven by the same structures discussed in the video game case. Increased advertising expenditures, falling cost and prices, and word of mouth communication all were driving sales. As Porter states,

> "Casual users were attracted to the market as a result of the energy crisis and the social trends back to nature, but the rapid growth of this segment had been accelerated by the introduction of many saws at rock bottom prices and the aggressive marketing of chain saws." (Porter 1981)

Several explanations for the crash stress factors that were external to the market. One such explanation contends that chain saw sales were driven by the sales of wood burning stoves and that the chain saw market peaked when stove sales began to decline. This explanation is questionable as there were several years in which chain saws sales and stove sales moved in the opposite direction. A second explanation holds that the

recession in 1980-1981 was responsible for the decline. This explanation is also questionable as chain saw sales never recovered significantly after the recession ended,

The firms in the chain saw market employed very different strategies. McCulloch and Beaird-Poulan both bet heavily on the casual users market and substantially expanded capacity by building new plants (Porter 1981, Business Week 1982). In the case of Beaird Poulan, their emphasis on the causal user market enabled them to increase their share of the overall market from 8% in 1972 to about 22% in 1977 while substantially increasing net income (Porter 1981).

Homelite adopted a more conservative strategy (Business Week 1982). Instead of building new plants, Homelite expanded incrementally (Business Week 1982, Porter 1981). Slow expansion produced the largest backlogs in the industry and Homelite's market share fell by 1% per year from 1973 until 1978 (Porter 1981). Homelite maintained its strong position in the professional market segment by continuing to improve its high-end saws and developing its dealer network.

The different firms' relative fortunes changed dramatically after the market peaked. McCulloch and Beaird-Poulan both suffered substantial losses due to excess capacity (Business Week 1982, Porter 1981). Black & Decker, McCulloch's parent company, finally sold McCulloch at a substantial loss in 1983 (Business Week 1982).

On the other hand, Homelite remained profitable even though its net income declined after the market peaked (Business Week 1982, Porter 1981). Homelite benefited from its conservative response to the casual user market and also gained some sales as some casual users traded up to more powerful saws (Porter 1981). By 1981, Homelite's share increased by 3%.

Homelite benefited from its conservative strategy in a growing market. A common prescription in growth markets is to expand capacity faster than the market growth rate in order to build market share. If a learning curve is present, a larger share leads to lower

costs and higher profit than the competition (Lilien 1992). Homelite followed the opposite strategy by expanding capacity slowly and losing share during the growth phase. Once the market collapsed, Homelite avoided excess capacity, maintained profitability and was positioned to dominate its nearly bankrupt competitors.

III. Conclusion

This chapter has demonstrated that firms across a wide range of industries have a great deal of difficulty managing rapid growth. Under- investment can cause firms to lose sales and market share during the growth and then suffer near financial collapse as the market slows and declines.

The next chapter presents a product life-cycle model that generates behavior patterns that are qualitatively similar to those observed in the case studies. The model served as the setting for the experiments reported in chapter four.

Chapter Four
The Simulation Model

This chapter presents the market simulation model. Section one presents the structure of the model while section two presents some illustrative simulations.

I. Model Structure

The experiments discussed in subsequent chapters are based on a computer simulation model of a rapidly growing market. The simulation model can be divided into three sectors: the market, the firm, and the competitor. Each of these sectors is described below.

The Market

The theoretical foundation of the market sector is the extensive literature dealing with product diffusion. In the marketing literature, a variety of models have been developed to explain the pattern of demand for a new product (Lilien 1992, Mahajan, Mueller & Bass 1990, Bayus 1993). While the details differ, these models include factors such as word of mouth and advertising that generate the growth of a new product. These models also include the penetration of the potential customer base that causes markets to saturate and decline.

Previous research has provided considerable support for product diffusion models. For example, Bayus examined sales data for 27 durable goods ranging from caluclators to waffle irons (Bayus 1993). He found that the Bass diffusion model provides an excellent fit to sales data. The model discussed below is typical of this general class of diffusion models.

In the simulated market, the subject's firm competes against a single competitor. The competitor can be interpreted as an individual competitor, in which case the market is a duopoly, or as an aggregate representation of several competitors.

The market sector is divided into two parts that determine "Total Market Orders" and "Market Share" respectively. "Total Market Orders" are the total number of orders received by both the firm and the competitor. "Market Share" is the fraction of "Total Market Orders" received by the firm. The firm's sales equals "Total Market Orders" multiplied by "Market Share".

The equation for "Total Market Orders" is:

(1) Total Market Orders = (Potential Buyers - Out of the Market-
Total Backlog)*Order Fraction
(2) Potential Buyers = 150E6*EXP(-(Reference Price*.2)^.15)
(3) Reference Price = Minimum(Price, Competitor Price)

Starting from the first term in equation 1, it is assumed that there is a certain number of households that would be interested in buying the product (Potential Buyers). The number of households willing to buy the product is determined by the minimum of the firm's and the competitor's price. For example, if price was $80, twenty million households might be willing to buy. If price was reduced to $40, perhaps thirty million households would be willing to buy.

The demand curve is shown in equation 2 and graphed in figure 1. The specific shape of the curve was drawn to be consistent with past studies of market demand (Lilien and Kalish 1986). Overall market demand is assumed to be price inelastic with an elasticity of .24 at the initial market price level of $80. The price elasticity becomes slightly smaller as market price declines.

Figure 1

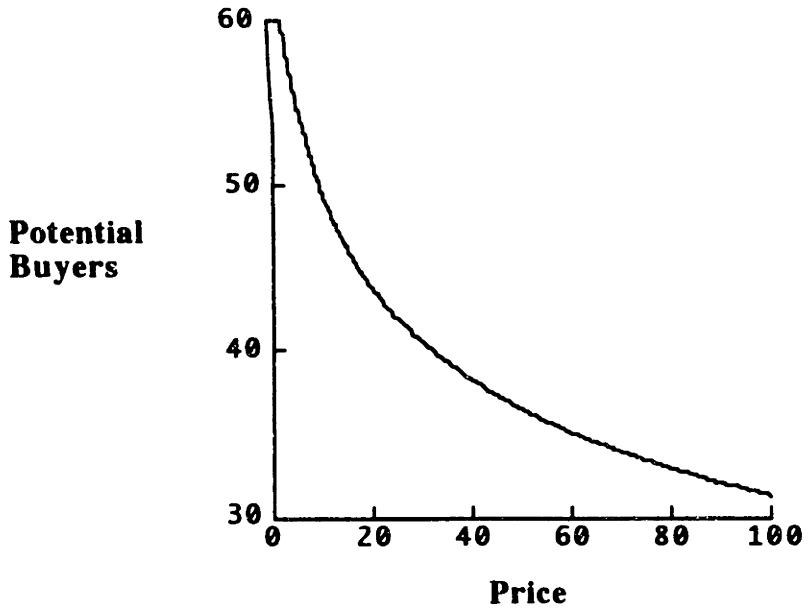

(4) Out of the Market(t) = Out of the Market(t-1) + (Total Sales - Reentering)

(5) Total Sales = Firm Sales + Competitor Sales

(6) Reentering = Out of the Market*Reenter Parameter

The variable "Out of the Market" is the number of people who have bought the product recently and will not buy it again for some period of time. Typically, in markets for durable goods, there is a significant delay between consumer purchases. For example, once someone buys a personal computer he is unlikely to buy a new one for several years. Each consumer is assumed to buy one unit of the product and then wait several quarters until he buys again.

The formulas that determine "Out of the Market" are given in equations 5 and 6. The inflow into "Out of the Market" is "Total Sales" which is the sum of the firm's sales and the competitor sales. The number of households that reenter the market is a fraction of

those that are currently "Out of the Market". The fraction, "Reenter Parameter", is a treatment variable that was chosen to generate various patterns of market demand.

(7) Total Backlog = Firm Backlog + Competitor Backlog

The variable "Total Backlog" is the total number of buyers who have ordered the product but have not yet received it. These individuals are not immediately eligible to purchase the product.

(8) Order Fraction = (Marketing Parameter*Total Effective Marketing Spending + WOM Parameter*Word of Mouth)/(1+ Marketing Parameter*Total Effective Marketing Spending + WOM Parameter*Word of Mouth)

(9) Total Effective Marketing = Firm Effective Marketing + Competitor Effective Marketing

(10) Word of Mouth(t) = Word of Mouth(t-1) + (Total Sales - Word of Mouth Decay)

(11) Word of Mouth Decay = Word of Mouth*.1

The "Order Fraction" controls the timing of total market demand by determining what percentage of the potential buyers purchase the product each quarter. The functional form of equation 8 is taken directly from a study of durable goods sales by Lilien and Kalish (Lilien and Kalish 1986).

The "Order Fraction" is determined by two factors. First, word-of-mouth communication increases the purchase fraction. As is common in the marketing literature, it is assumed that recent buyers influence non-buyers to purchase the product (Lilien 1992). In short, recent sales lead to more sales. The strength of the word of mouth process in influencing potential buyers is determined by the "WOM Parameter".

As will be discussed in Chapter 4, the "WOM Parameter" is a treatment variable that is varied experimentally to generate different market scenarios.

Further, it is assumed that the word of mouth generated by recent sales decays over time. This represents the assimilation process whereby people eventually take their purchases for granted. The decay process is represented by the variable "Word of Mouth Decay". The decay rate is 10% per quarter which is consistent with the value estimated in previous studies of durable goods (Lilien and Kalish 1986).

Second, total marketing spending, by both the firm and the competitor, increases the "Order Fraction". Marketing spending provides consumers with information about the product and may give them an incentive to buy. The power of marketing spending to induce purchases is determined by the "Marketing Parameter". The value of the "Marketing Parameter" was chosen to produce reasonable patterns of market demand when combined with the "WOM Parameter" treatment variable.

(12) Firm Effective Marketing(t) = Firm Effective Marketing(t-1) + (Firm Marketing - Firm Effective Marketing(t-1))*.2

Previous studies have shown that there is a delay before a change in marketing spending has its maximum effect (Lilien & Kotler 1992). This delay represents the total time required to plan and implement marketing programs as well as the time required for a message to "sink in". In equation 12, "Firm Effective Marketing" is an exponential average of past levels of marketing spending. The exponential average places more weight on recent value of marketing spending. The adjustment coefficient of .2 implies that 67% of the impact of a decision to change marketing spending is realized within 5 quarters.

The relationships that control total market orders form a set of feedback processes that are collectively referred to as the growth-saturation loops. As shown in figure 2, an increase in sales generates word of mouth communications which further increases orders and sales. Further, an increase in sales increases revenue and market spending which generates still more orders and sales. These two processes are self-reinforcing and produce rapid growth of sales. Growth is limited by the market saturation loop. Increased sales remove potential buyers from the market. An increase in buyers who are out of the market reduces total market orders.

Figure 2

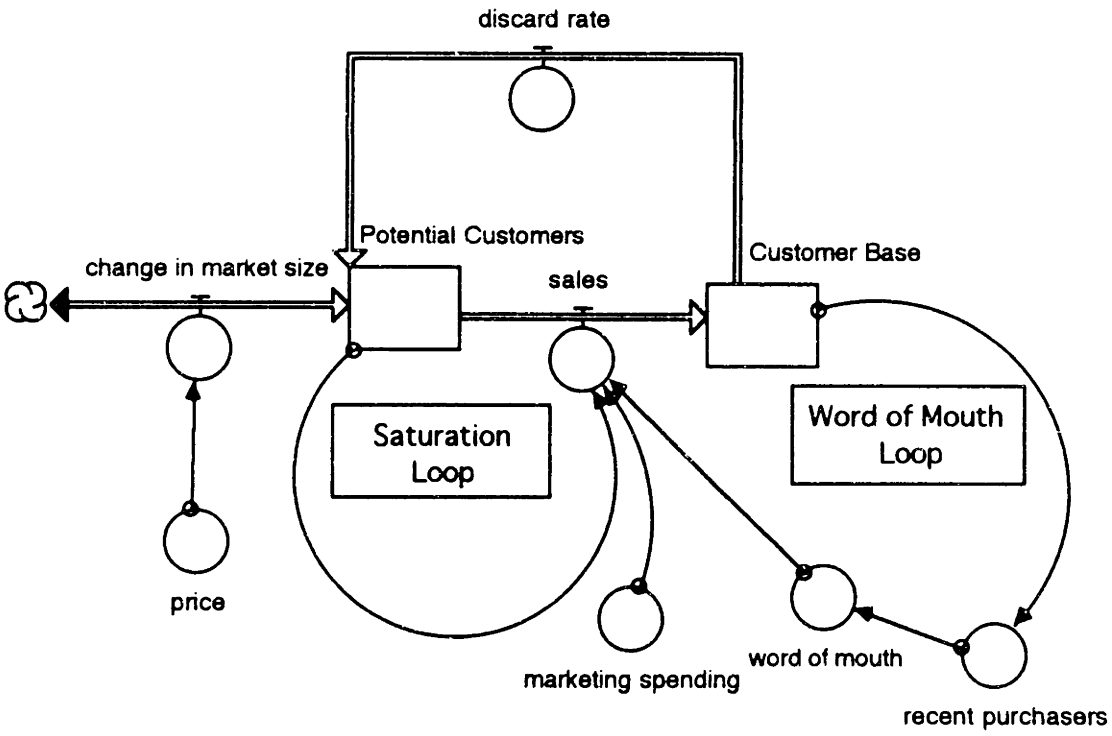

Market Share

(13) Market Share = Firm Attractiveness/(Firm Attractiveness + Competitor Attractiveness)

(14) Firm Attractiveness = Attractiveness Constant*(Firm Effective Marketing^a1)*(Firm Price ^ a2)* Firm Delivery Delay(^a3)

(15) Competitor Attractiveness = Attractiveness Constant*(Competitor Effective Marketing^a1)*(Competitor Price ^ a2)* Competitor Delivery Delay(^a3)

(16) Delivery Delay = Backlog/Sales

In the model, market share is determined by the attractiveness of the firm's product relative to the competition (Lilien & Kotler 1992). Attractiveness is related to marketing spending, price, and delivery delay. It is assumed that higher marketing spending, lower price, and a shorter delivery delay all increase attractiveness.

The equation for "Market Share" uses a functional form that is often referred to as US/(US+THEM) (Lilien & Kotler1992). "Market Share" is equal to the firm's attractiveness divided by the sum of the firm's attractiveness and the competitor's attractiveness. This calculation must result in a number between zero and one. This functional form assumes that an increase in a component of attractiveness produces ever smaller increases in market share as market share becomes larger.

The equation for attractiveness uses a multiplicative functional form that is common in the marketing literature (Lilien & Kotler1992). The multiplicative attractiveness function combined with the US/US+THEM market share implies a simple formula for the elasticity of market share with respect to the components of attractiveness. For example, the price elasticity of market share is,

(16) e(price) = a2*(1-Market Share)

The market was assumed to be highly price competitive. The coefficient a2 was assumed to be -5 which, given an initial market share of 50%, implies a price elasticity of market share equal to -2.5. This elasticity is reasonable for a highly competitive consumer goods markets (Tellis 1988).

Previous studies have usually found diminishing returns to marketing spending which implies that the effectiveness coefficient for marketing spending should be less than one (Lilien & Kotler 1992). In addition, a meta-analysis of 41 previous studies covering 367 price elasticities for 220 brands, demonstrated that the price elasticity is often 8-10 times larger than the advertising elasticity (Tellis 1988). After experimentation with the simulation model, a value for a1 of .4 was found to generate reasonable competitive dynamics. With an initial value of market share of -50%, the elasticity of market share with respect to marketing spending is .2 compared with a price elasticity of 2.5.

Product availability has been shown to have a significant impact on sales (Stalk & Hout 1990). However, in contrast to marketing spending and price, there are virtually no studies that have formally estimated the elasticity of demand with respect to availability. The parameter a3 was set equal to -1 which implies an elasticity of -.5. This value implies that changes in lead time are less important than price changes but more important than changes in marketing spending.

The firm's orders are affected by the size of its backlog relative to its shipments. If the backlog is very large, relative to shipments, the firm will have a long delivery delay. For example, if backlog is 100 units and shipments are 50 per quarter, it will take 2 quarters before a new order can be shipped.

The Firm

The simulation deals with the evolution of a new market from the growth stage through the maturity stage. The player's firm sells a durable product in a market with one competitor. Because the product was introduced recently, the firm has had only two quarters of sales experience. The competitor introduced a similar product at about the same time.

(17) Firm Backlog(t) = Firm Backlog(t-1)+ (orders-sales-cancellations)

(18) Orders = Potential Orders*Market Share

(19) Cancellations = Backlog*.10

(20) Sales = Minimum(Backlog - Cancellations, Capacity)

The player's firm receives orders for its product which accumulate in an order backlog until they can be shipped. The firm will sell its entire order backlog (less any cancellations) if it has adequate production capacity. Otherwise, it will sell an amount equal to its production capacity. For example, if the firm has 20,000 units of capacity and a backlog of 10,000 the firm will ship 10,000 units. If the firm has a backlog of 30,000 units, it will ship 20,000 units.

Each quarter, the firm loses 10% of its orders due to cancellations. The cancellations take place because people get tired of waiting for their merchandise. The flows in and out of the firms backlog are shown in figure 3.

Figure 3

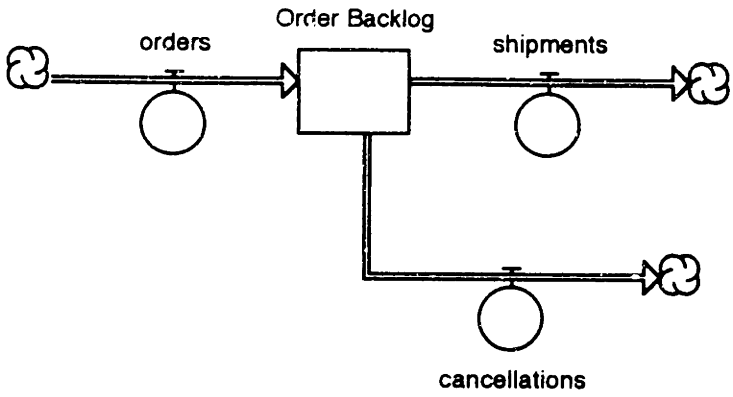

(21) Capacity(t) = Capacity(t-1) + Capacity Change

(22) Capacity Change = (Target Capacity - Capacity)/4

The player controls profitability by manipulating target capacity and price. Target Capacity is the firm's desired level of productive capacity measured in units per quarter. At the beginning of a quarter, the firm has a stock of plant and equipment available for production. The stock of plant and equipment translates into the maximum amount that the firm can produce and ship during a quarter. The firm can never produce more than its capacity. To simplify the game, it is not possible for the firm to build and hold inventory.

It takes time for the firm to adjust capacity to its target level. In the simulation, it takes an average of four quarters to acquire a unit of capacity. However, the four quarters is an average; some capacity can be added sooner while other expansion projects take longer.

Figure 4 shows how capacity responds to a one-time increase in Target Capacity. At quarter 5, Target Capacity increases from 1,000 units to 2,000 units per quarter. Capacity immediately starts to grow as the firm does the easiest expansion projects first.

Capacity gradually approaches Target Capacity as the firm completes longer, more difficult projects.

Figure 4

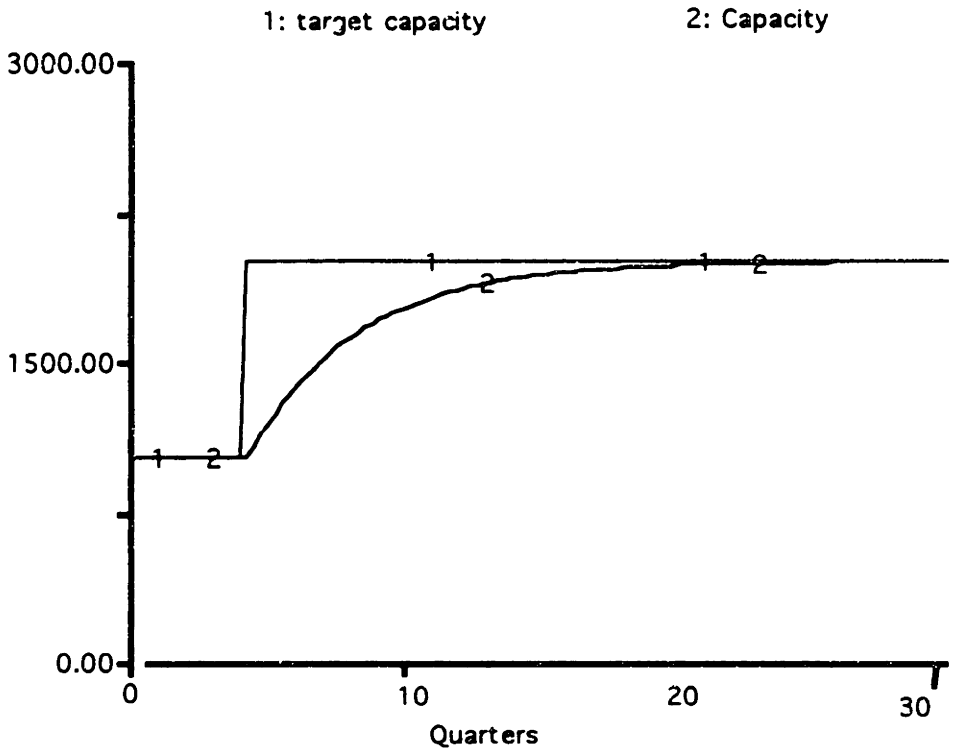

Price is the second decision variable. As discussed above, a lower price increases both total market demand and the firm's market share.

(23) Marketing Spending = RevenueL1*.05

(24) RevenueL1 = Delay(Revenue,1)

Marketing Spending is the total amount the firm spends on advertising and sales promotion measured in dollars per quarter which, as explained below, increases total market demand and the firm's market share. The firm's marketing spending is set at 5% of the last quarter's revenue. The variable "RevenueL1" is revenue lagged one quarter.

The components of the firm's revenue and costs are shown are as follows.

(25) Revenue = Firm Price*Firm Sales

"Revenue" is the product of of the firm's price and unit sales.

(26) Variable Cost = Sales*Unit Variable Cost

(27) Unit Variable Cost = 30*(Cumulative Sales/Initial Cumulative Sales)^-.3

"Variable Cost" is determined by orders and unit variable cost. Unit variable cost starts at $30 per unit and falls as the firm's cumulative sales increase. The reduction in unit variable cost caused by higher cumulative sales simulates the behavior of a learning curve. Unit variable cost is assumed to fall 20% with every doubling of cumulative sales (Lilien & Kotler 1992)

(28) Fixed Cost = Capacity*30

"Fixed Cost" is determined by capacity and unit fixed cost. It is assumed that each unit of capacity generates debt service and overhead cost each quarter. Unit fixed cost is constant at $30 per unit of capacity.

(29) Investment Cost = Absolute Value(Change in Capacity)*45

"Investment Cost" is what the firm must spend when it increases or decreases capacity. These costs include installation costs and the cost of scrapping capacity. Investment cost is determined by the absolute value of the change in capacity during a

quarter and the unit adjustment costs. The absolute value is taken so that reductions in capacity genrate positive investment costs. The unit adjustment cost is constant at $45.

(30) Net Income = Revenue - Variable Cost - Fixed Cost - Marketing Spending - Investment Cost

"Net Income" is the difference between revenue and the sum of all the costs.

Variables such as product quality and functionality are assumed constant. These assumptions are made in order to simplify the structure of the game. The firm can influence sales only by controlling price and target capacity.

Players must base their decisions on data provided by the firm's information system. Every quarter, the player receives a report about how the previous round of decisions affected sales, costs, and profitability. The report enables the player to examine all previous information about past decisions and the firm's performance.

The Competitor

The structure of the competitor is identical to that of the firm. The only difference is that the competitor's decisions for price and target capacity are determined by decision rules instead of by a player.

(31) Competitor Target Capacity = Competitor Orders L1 + 4*Competitor Expected Growth + Competitor Backlog Gap

(32) Competitor Expected Growth = Competitor Orders L1*Average Competitor Growth Rate

(33) Average Competitor Growth Rate = Smooth(Competitor Orders,2)

(34) Competitor Growth Rate = (Competitor Orders L1 - Competitor Orders L2)/Competitor Orders L2

(35) Competitor Orders L1 = Delay(Competitor Orders,1)

(36) Competitor Orders L2 = Delay(Competitor Orders,2)

(37) Competitor Backlog Gap = Competitor Backlog - Competitor Production Capacity

The competitor's decision rule for target capacity based on last quarter's orders ("Competitor Orders") , an average of past growth rates ("Competitor Expected Growth") and the size of the order backlog ("Competitor Backlog Gap"). The competitor decides target capacity by extrapolating an average of past order growth rates four quarters into the future. In addition, target capacity is increased if the competitor's order backlog is large relative to production capacity.

"Competitor Expected Growth" is an exponential average of past growth rates with an averaging time of two quarters. The "Competitor Growth Rate" is the percentage growth rate of orders. "Competitor Orders L1" and "Competitor Orders L2" are the competitor's orders lagged one and two quarters respectively.

"Competitor Backlog Gap" is the difference between the competitor's order backlog and the competitor's production capacity. The competitor increases (decreases) target capacity if backlog if greater (less) than production capacity.

(38) Competitor Price = (30 + Competitor Unit Variable Cost)/(1-Competitor Target Margin)

The competitor's decision rule for "Competitor Price" is based on unit variable cost and a constant markup percentage. Unit cost is the sum of variable production cost and the fixed cost ($30) of producing one unit of product. The markup percentage is 30%.

The competitor's price falls as unit variable cost declines according to the learning curve. The learning curve for the competitor is the same as that for the firm.

II. Model Simulations

Word of Mouth and Purchase Frequency

The purpose of the first set of simulations is to explore the dynamics of the market sector in isolation. In the market sector, marketing spending is held constant at $50,000 per quarter. Growth is generated by the positive feedback loop created by word of mouth communication. Demand is limited by the market saturation loop.

The strength of the word of mouth loop can be varied by changing a single parameter (WOM). A larger value of WOM means that each recent purchase leads to more additional sales.

The strength of the market saturation can be varied by changing the fraction of potential buyers who reenter the market each quarter (RENT). A large (small) value of RENT means that buyers repurchase the product frequently (infrequently).

The size of the WOM parameter influences the height of the peak in sales. Figure 5 shows two simulations of the market model with WOM equal to its base value and twice the base value. The larger the value of WOM, the faster the growth and subsequent decline in sales. With the larger WOM, the overshoot is 31% greater than in the base case.

Figure 5

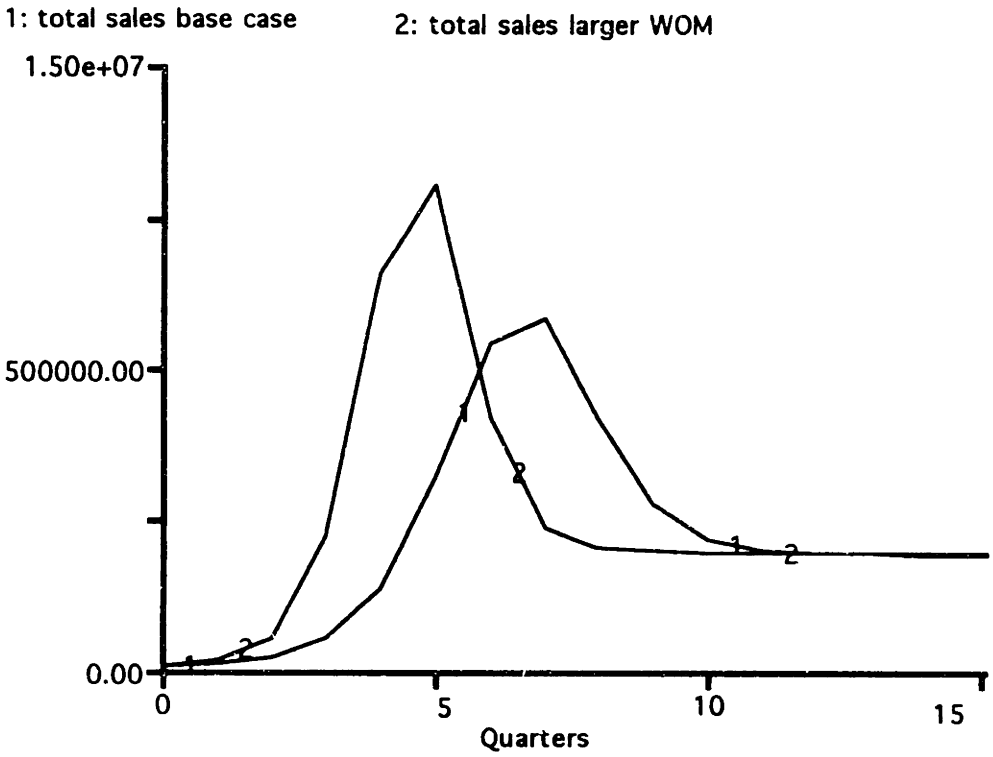

The size of the RENT parameter influences the size of the decline in sales following the peak. Figure 6 compares the base case to a simulation in which RENT is twice the base level. In both simulations, sales peak and then fall to the long run steady-state level. The larger the value of RENT the more often people buy the product and the larger the steady-state level of sales. Sales don't have very far to fall when RENT is large. Overall, the magnitude of the overshoot declines by 43% with the doubling of RENT.

Figure 6

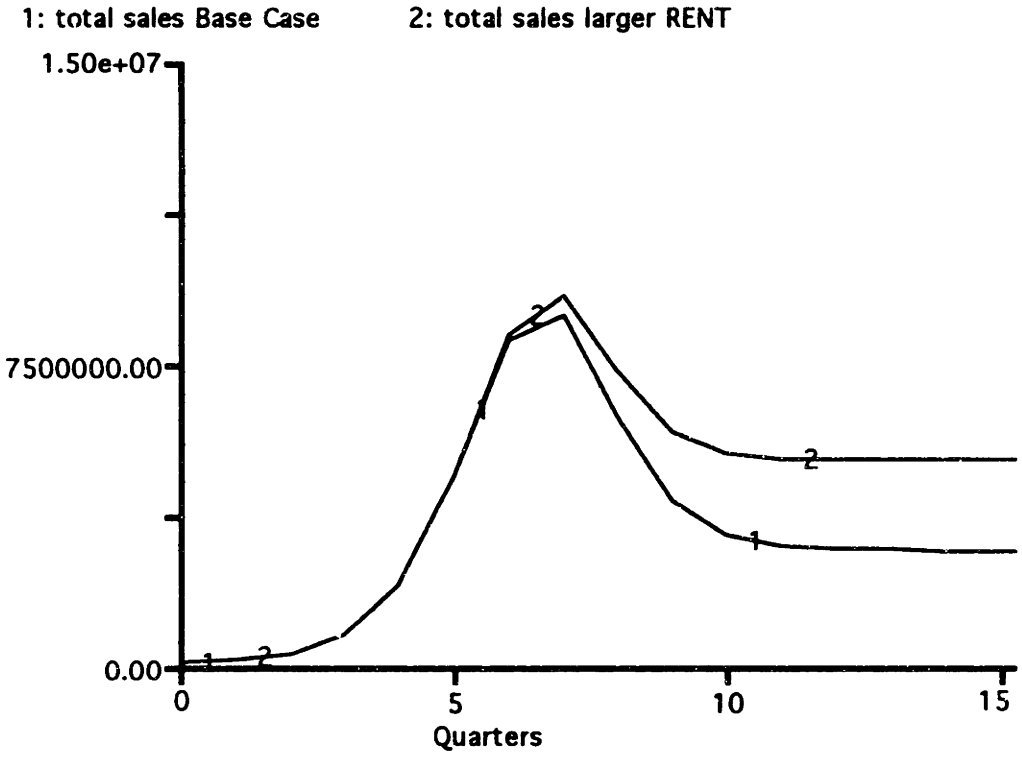

Learning Curve

In the third pair of simulations, learning curve dynamics are added to the market sector. In these simulations it is assumed that unit variable cost falls by twenty percent with each doubling of cumulative sales. Price is equal to the sum of unit variable cost and unit fixed cost increased by a constant margin. Price falls as unit variable cost declines. The equation for price is identical to the competitor price equation used in the full simulation (equation 38).

Figure 7 compares two simulations. In the "base case simulation", the learning effect is neutralized and unit variable cost is constant at $30. In the second case, unit variable cost falls 30% with every doubling of cumulative sales. Adding the learning effect causes sales to peak at a higher level and to stabilize at a slightly higher equilibrium. Overall, the magnitude of the overshoot increases by about 4% when learning is included.

Figure 7

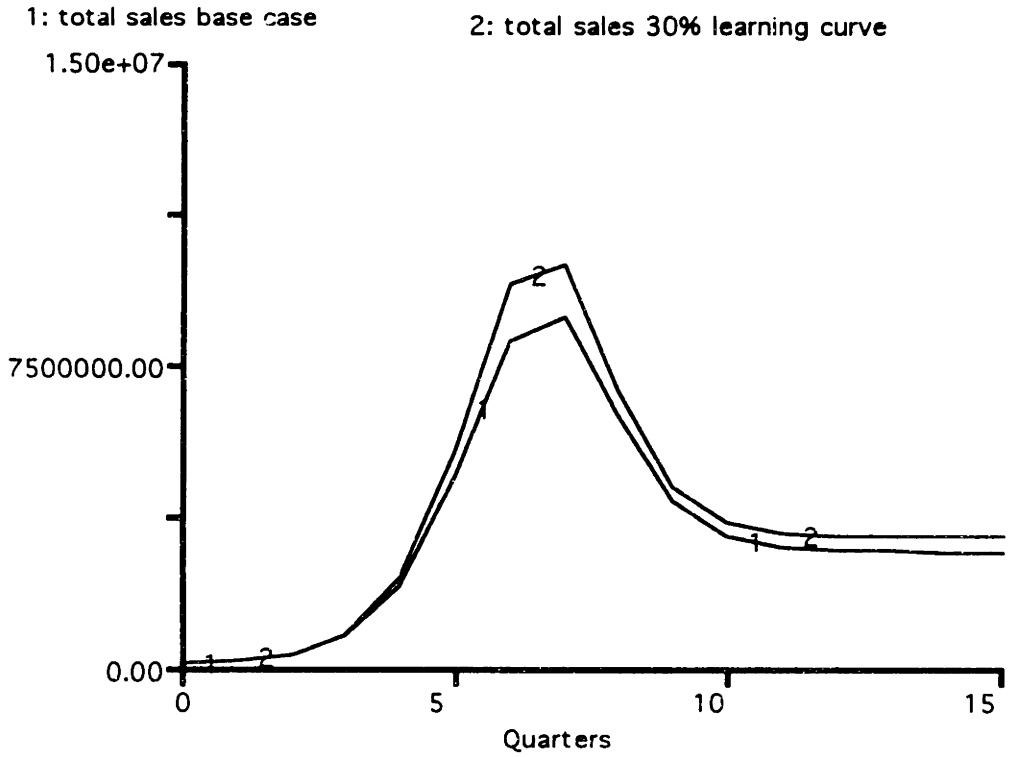

Variable Marketing Spending

In the fourth pair of simulations, marketing spending varies with total sales and price. In these simulations, it is assumed that firms spend eight per cent of revenue

(price*sales) on marketing. This simulates a situation in which firms invest more in marketing as their revenues grow. In addition, it takes an average lag of five quarters before increased marketing spending affects sales.

Figure 8 compares two simulations. In the "base case simulation", marketing spending is constant at $100,000. In the second case, marketing spending varies as described above. Allowing marketing spending to vary causes sales to peak at a higher level and to stabilize at a somewhat greater value. Overall, the magnitude of the overshoot increases by about 10% when variable marketing spending is added.

Figure 8

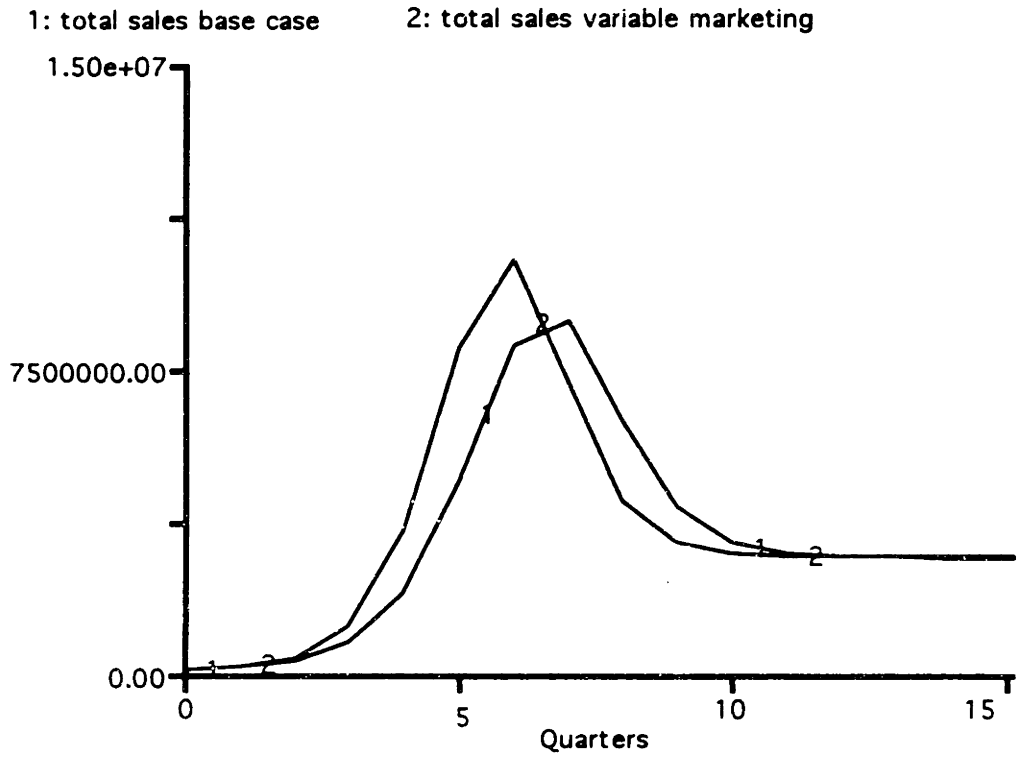

Capacity Management

This set of simulations uses the complete model described at the beginning of this chapter. In the model, orders and sales are constrained by the availability of production capacity. Market share is lower if the order backlog is large relative to production capacity. In addition, the firm can't sell any more than it can produce.

The simulated firm's decision rules for target capacity and price are identical to those of the competition . As explained in the section on the competitor's decision rules, the inputs into the capacity decision are last quarter's orders, an average of past growth rates and the size of the order backlog. The firm sets target capacity by extrapolating an average of past order growth rates four quarters into the future. In addition, target capacity is increased if the firm's order backlog is large relative to production capacity.

The firm's decision rule for price is based on unit cost and a constant markup percentage. Unit cost is the sum of variable production cost and the fixed cost of producing one unit of product. The markup percentage is 30%. The firm's price falls as unit variable cost declines according to the learning curve.

Figure 9 shows the model behavior of target capacity and capacity. Target capacity and increase rapidly during the market's growth phase. Eventually, the market peaks and both capacity and target capacity decline.

Figure 9

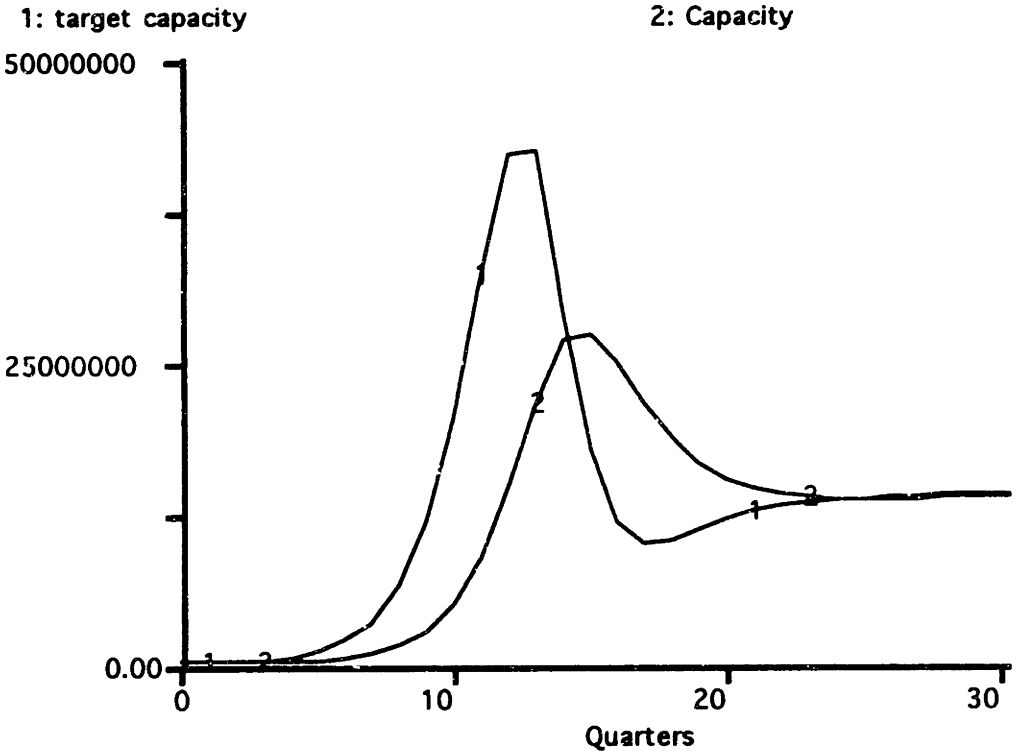

Capacity doesn't fall as fast as sales. The sales decline is caused by the growth-saturation feedback loops described above. These processes are capable of generating a very rapid sales decline. The firm doesn't anticipate the decline and there is a delay before it can reduce capacity. Capacity utilization, sales and net income are shown in figure 10. The excess capacity continues to generate costs and net income declines sharply. Eventually, the firm eliminates the excess capacity and restores profitability.

Figure 10

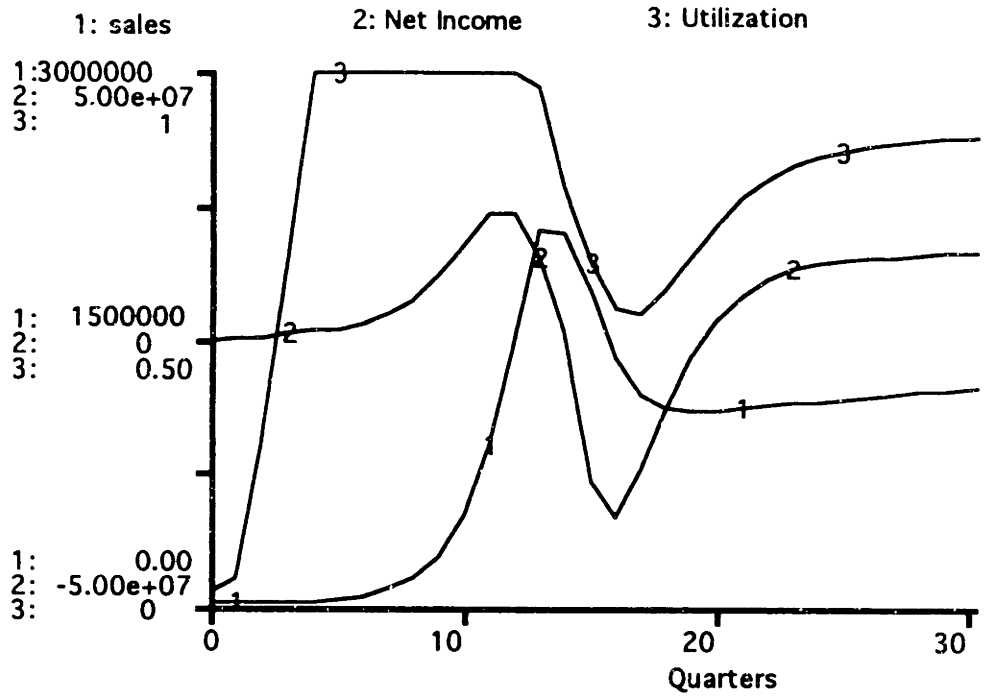

Figure 11 shows the behavior of price. Early in the simulation, price declines rapidly as the learning curve drives down unit variable cost. The price cuts become progressively smaller as unit variable cost becomes very small.

Figure 11

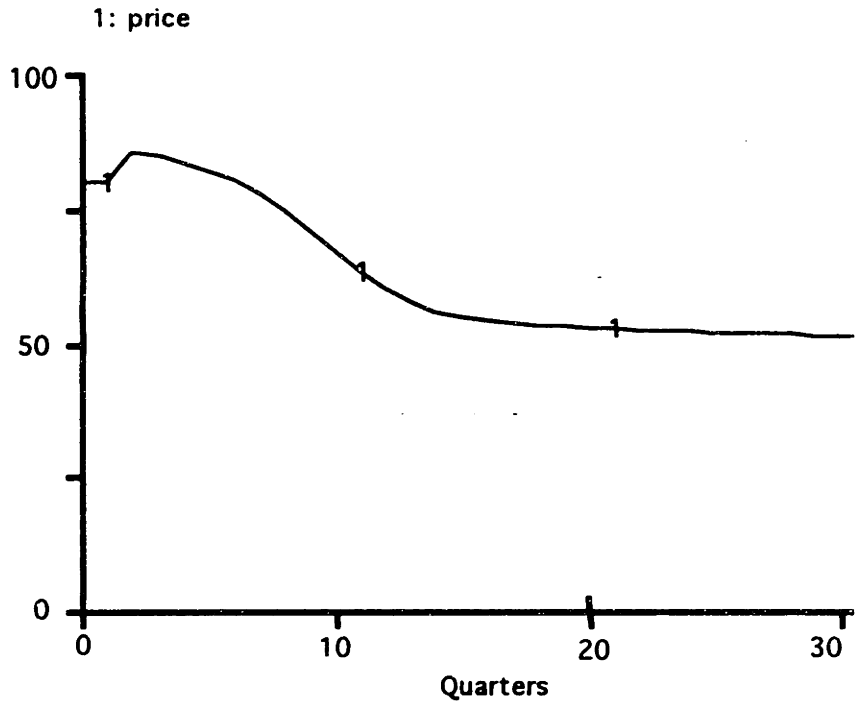

Qualitatively, the simulation is very similar to the case studies described in chapter two. The market sector of the model generates very rapid growth and collapse without the effect of external factors such as recessions or changes in interest rates. In the base case, sales peak at quarter thirteen and then decline 55% over the next eight quarters. The magnitude of the decline is similar to what was experienced in video games and home computers. The case studies showed that firms often don't anticipate the decline in sales which causes substantial excess capacity and extreme declines in profit. Again, the behavior pattern generated by the model is qualitatively similar.

Chapter Five
Experimental Design

This section describes the experiment in dynamic decision making. The first section reviews the structure of the task described in the previous chapter. The second section discusses the hypotheses and the design of the experiment. The third section presents results.

I. Task

Chapter four described the structure of the growth market model and explained how the model generated market growth and collapse. In the experiment, the model was converted into an interactive simulation. The objective of the game was to maximize cumulative net income over forty quarters. The determinants of revenue and cost in the game are the same as those described in chapter four. The simulation game was identical to the model with the exception that price and target capacity were determined by the subjects instead of by decision rules.

The simulation included an artificial firm that was managed by the computer. Decisions made by the artificial firm were determined by a set of simple rules The artificial firm in the interactive simulation used the same decision rules as in the model described in chapter four.

The subjects decided on target capacity and price in each of 38 quarters. Target capacity was the amount of capacity the firm would like to have. Actual capacity gradually adjusted toward target capacity with a four quarter delay. The speed of adjustment was fixed in the model which forced the subjects to anticipate changes in demand. Price was simply the price per unit of product.

The subjects took control over the firm in quarter three. Target capacity and capacity were initially both equal to 20,000 units and price was equal to $80 per unit. The results for the first two quarters were determined by the initial conditions of the simulation and were displayed on the spreadsheet when the subjects opened the game file. The initial conditions for the game were representative of the introduction phase of a new product. Sales were low but growing rapidly. Net income was negative because demand was much less than the initial 20,000 units of capacity, but capacity utilization and profit were both increasing due to rapid market growth. The competitor was expanding capacity in anticipation of future demand and lowering price as he moved down the learning curve.

II. Hypotheses and Experimental Design.

Sterman has argued that peoples' mental models do not adequately account for the effects of feedback loops, non-linearity, and time delays (Sterman 1989a,1989b). This "misperception of feedback structure" causes decision makers to use rules that are poorly matched to the dynamic environment. Decision rules that produce good performance in one setting can fail completely in an environment with a different feedback structure.

Previous research on industry evolution and the product life cycle has shown that two sets of feedback relationships are fundamental to market dynamics (Lilien & Kotler 1992). The first set of loops are the self-reinforcing processes that drive market growth. Two examples of growth loops are the dynamic processes that govern word of mouth and advertising discussed in chapter three. The second crucial feedback process is the negative loop that generates market saturation and decline. Chapter four explained how the progressive exhaustion of the stock of potential customers can cause a sudden decline in market demand.

The misperception of feedback hypothesis has significant implications for the subjects' expected performance in the life cycle simulation. The hypothesis suggests that relative performance might vary with the intensity of the two key loops. For example, many texts on marketing or business strategy emphasize strategies that focus on the learning curve (Porter 1985). During the market's growth phase, businesses are directed to expand capacity in advance of market demand in order to capture the largest market share. In the presence of a strong learning curve, the largest market share produces the lowest cost, and the highest profit. This strategy, which is appropriate when market growth is smooth, could be disastrous when the word of mouth and saturation loops are powerful. Strengthening these loops causes more violent market overshoot and collapse and can produce huge losses due to excess capacity (Beinhocker, Henderson, Sterman 1992).

The misperception hypothesis also has implications for the nature of learning that would be expected in the simulation. The first type of learning that could have taken place is a general learning effect that occurred as the subjects became more familiar with the information display, the experimental apparatus, and, in later trials, the basic magnitude and pattern of market demand. The second type of learning is the acquisition of robust mental models that minimize the impact of changes in feedback loop intensity. It would not be surprising if the subjects had difficulty managing the more difficult scenarios on the first trial, but the efficacy of their strategies should improve with experience The misperception hypothesis suggests that learning robust decision rules is much more difficult than general learning and may not have occurred at all.

To test the above hypotheses, five different scenarios were created by varying two of the market parameters: the strength of word of mouth communication (WOM) and the average time between purchases (RENT). As shown in table 1, the design was a two by two with one control scenario. Table 1 shows the parameter values associated with

scenarios two through four. In the four scenarios, the word of mouth parameter and the average time between purchases were set to one-half and twice the level of the control scenario.

	Table 1	
	WOM	
RENT	2	.5
2	2	3
.5	4	5

The word of mouth parameter (WOM) controlled the strength of a positive feedback loop that generates market growth. As explained in chapter 4, a larger value of WOM caused demand to grow faster and to peak at a higher level. Given a fixed base of potential customers, a higher peak caused a larger overshoot of market demand above its equilibrium level. The fraction of customers reentering the market (RENT) controlled the strength of a negative feedback loop that limited market growth. A small value of RENT meant that there was a long delay between purchases and that the steady-state level of sales was relatively low. A low steady-state level of sales increased the overshoot of sales. The five scenarios were generated by values of WOM and RENT that were one-half and twice the size of the baseline parameters.

Each subject played five trials of the game with a trial consisting of 40 quarters. A latin square design was used in which the scenarios were arranged in five order patterns and subjects were randomly assigned to each pattern (Neter, Wasserman & Kutner 1985). The order patterns are shown in table two. For example, subjects assigned to the first order pattern played scenario one first and scenario two second. Subjects assigned to the second order pattern played scenario two first and scenario three second.

		Table 2		
Pattern 1	Pattern 2	Pattern 3	Pattern 4	Pattern 5
1	2	3	4	5
2	3	4	5	1
3	4	5	1	2
4	5	1	2	3
5	1	2	3	4

The sample consists of 122 students enrolled in two sections of an elective course in system dynamics. About 75% of the subjects were MBA students, 10% were mid-career middle managers in an executive education program, and the remainder were MIT undergraduates and graduate students from other MIT departments. Although a few subjects failed to complete or report all five trials, the experiment was reasonably well balanced. As shown in table 3, the number of trials in each of the twenty five cells (five trials by five scenarios) ranged from 19 to 27 with a mean of 23.

			Table 3			
			Scenario			
Trial	**1**	**2**	**3**	**4**	**5**	**Total**
1	21	25	21	19	23	109
2	25	24	27	22	21	119
3	21	22	22	24	23	112
4	22	21	23	23	24	113
5	24	23	22	22	21	112
Total	113	115	115	110	112	565

The task was assigned as a homework problem set to be completed in ten days. The subjects were allowed to spend unlimited time on the simulation during the ten day period and to suspend play between trials. The assignment instructed the subjects to keep a log during the trials which explained the strategy they intended to use and an evaluation of their results. The subjects received a grade on the write-up based on the quality of their analysis. In addition, the subjects received bonus points toward their class grade based on the cumulative net income earned over the five trials. Grades were not reduced regardless of how poorly a subject did in the game. The graded written assignment and the bonus point system provided a powerful incentive for the subjects to formulate intelligent strategies and, in general, to take the game seriously. The subjects were clearly concerned with grades as evidenced by the large number of questions about the bonus point system.

The information display instantaneously provided feedback about the resulting revenue, costs, and other market variables after a subject entered his capacity and price

decisions. Figure 1 shows the information display used in the game. During a trial, a subject could use the spreadsheet to look at the results from any previous quarter in that trial and plot any variable as a graph over time. The subjects could not access information from previous trials. The computer automatically recorded the subjects' decisions, all the feedback information, and the time required to play the game.

Figure 1

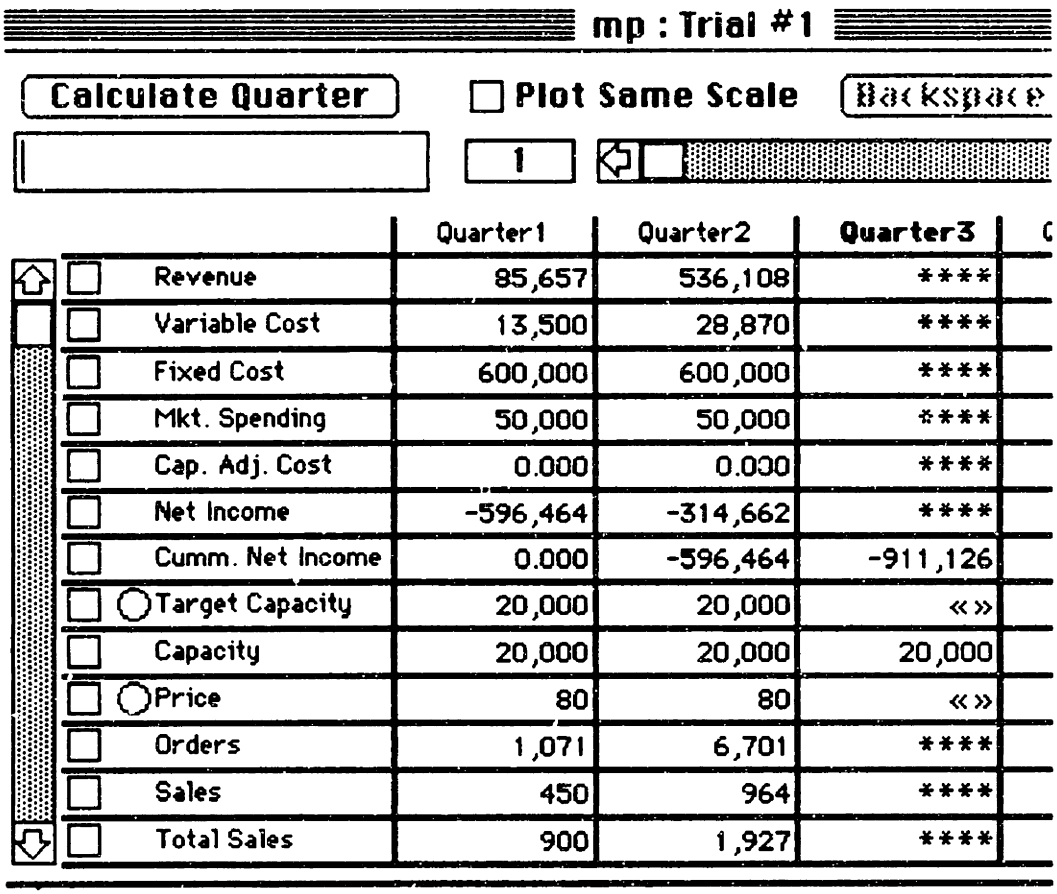

	Quarter1	Quarter2	Quarter3
Revenue	85,657	536,108	****
Variable Cost	13,500	28,870	****
Fixed Cost	600,000	600,000	****
Mkt. Spending	50,000	50,000	****
Cap. Adj. Cost	0.000	0.000	****
Net Income	-596,464	-314,662	****
Cumm. Net Income	0.000	-596,464	-911,126
Target Capacity	20,000	20,000	« »
Capacity	20,000	20,000	20,000
Price	80	80	« »
Orders	1,071	6,701	****
Sales	450	964	****
Total Sales	900	1,927	****

The structure of the market system and the details of the interactive simulation were presented in a class lecture and explained in written directions. The directions are included in appendix one of this chapter. The subjects were informed of the determinants of market demand and market share but were not given specific parameter values. The subjects were given no information about the competitor's decision rules or profitability.

III. Results

Behavior Patterns

Figures 2, 3 and 4 show typical sets of results. In figure 2, a subject playing his third trial rapidly expanded capacity in response to market growth. At quarter 15, market sales peaked and declined sharply. The subject reacted slowly to the market decline and was left with excess capacity and huge losses. Eventually, capacity and demand came into balance and profitability was restored.

In figure 3, a different subject playing her third trial followed a very conservative strategy. Capacity expanded very slowly which led to large backlogs and long delivery times. Consequently, market share declined continuously during the market's growth phase.

Figure 2

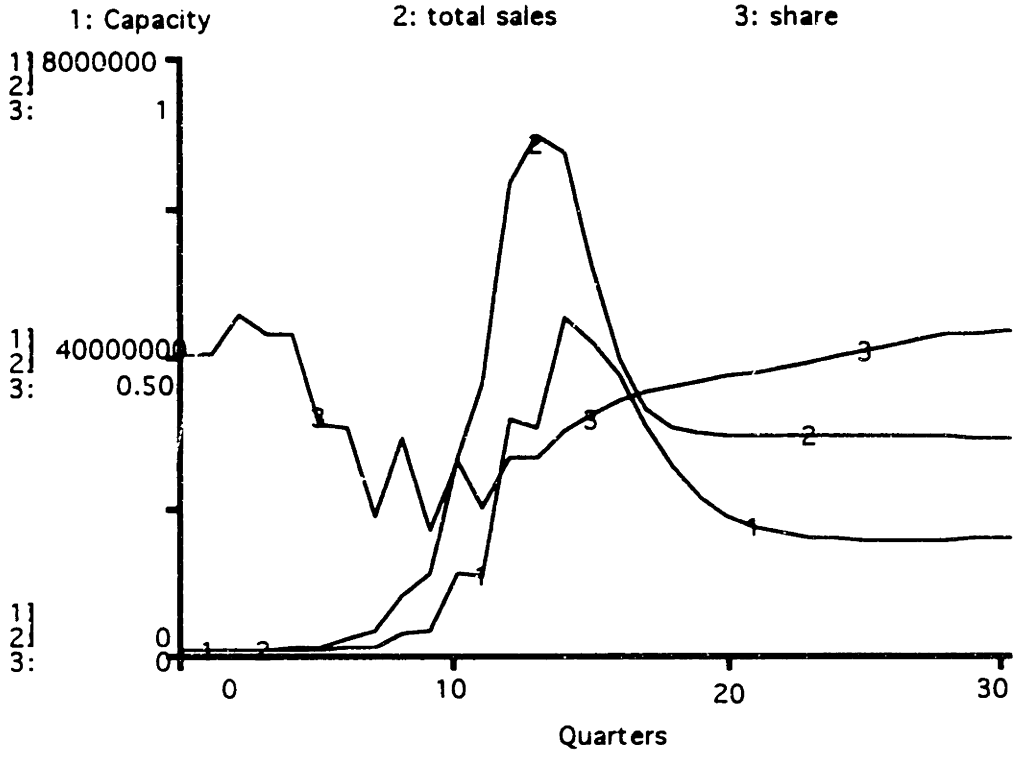
1: Capacity
2: total sales
3: share
8000000
1
4000000
0.50
0
0
0
10
20
30
Quarters

Figure 3

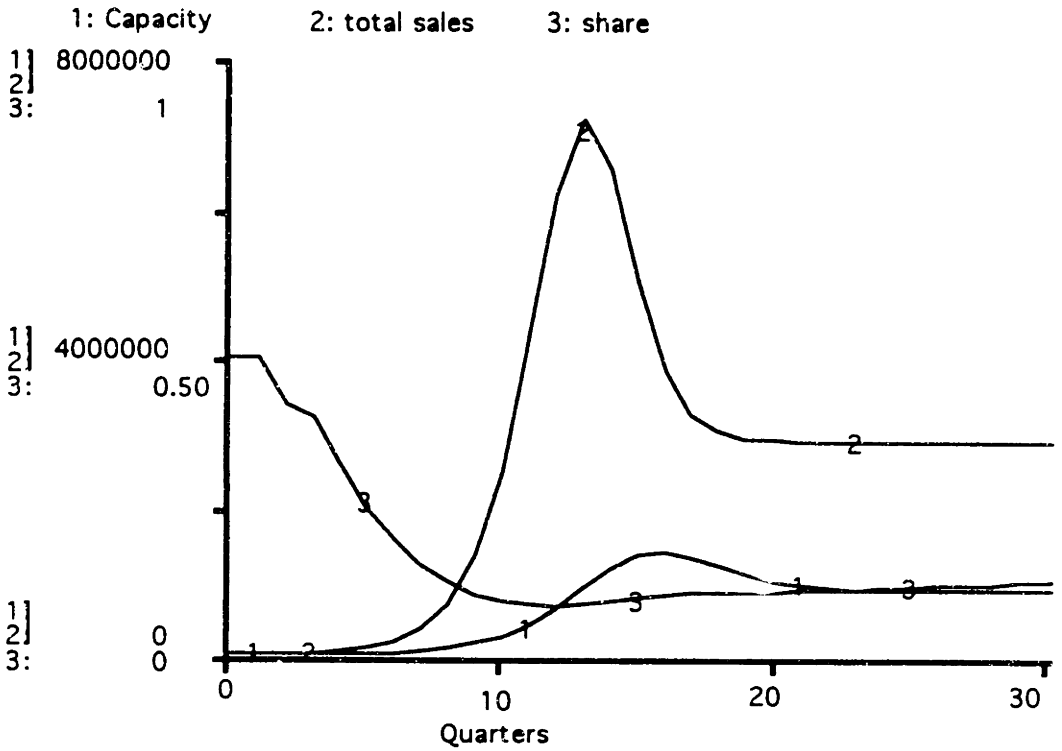

Once the market peaked, backlogs declined and market share improved slightly. Market share did not rebound to the original 50% partly because variable cost and price were higher than the competitor's. The higher cost was caused by the learning curve dynamics. The subject's firm did not move as far down the learning curve because its cumulative volume was less than the competitor's.

Figure 4 displays what can be described as a "ballistic" strategy. In a ballistic strategy, subjects set target capacity or price based on experience from previous games or other external considerations. The subject's decisions were not responsive to feedback about orders, market size, or growth from the current game. For example, in figure 4, a subject playing his fifth trial significantly increased capacity in advance of the market and then left it at the same level for most of the rest of the trial.

Figure 4

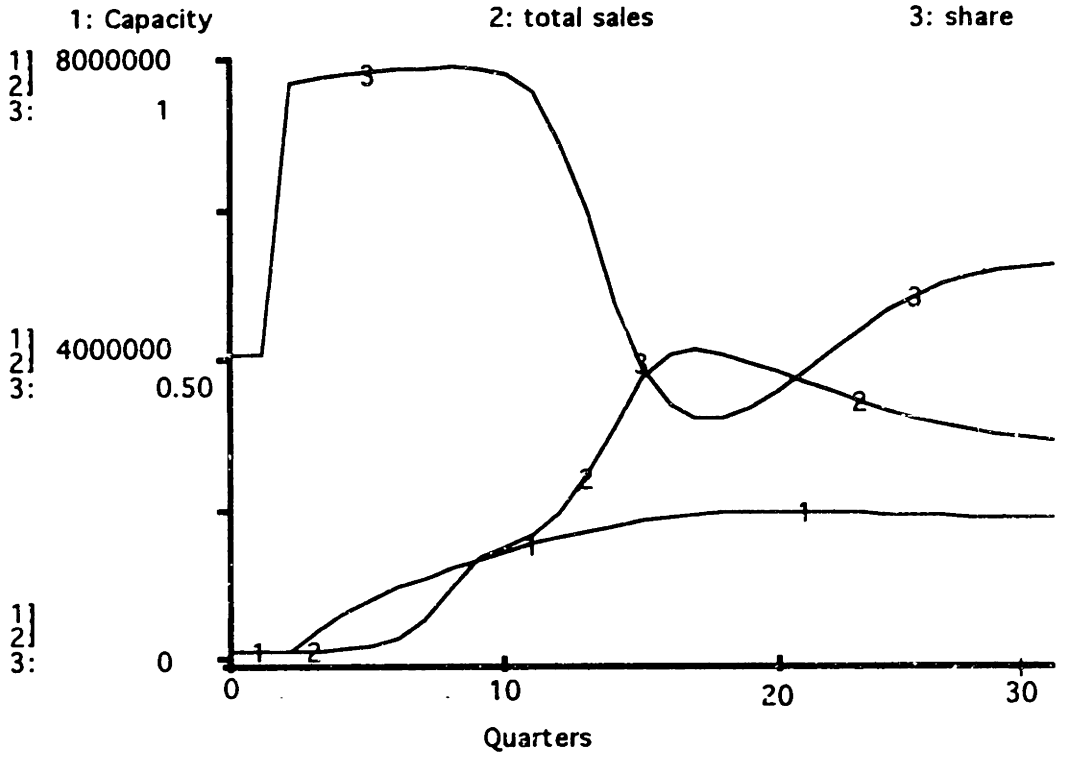

The differences between the three examples are illustrated by the causal loop diagram shown in figure 5. The diagram includes the positive feedback loops that generate growth and the negative market saturation loop that was discussed in chapter three. The diagram also includes the negative feedback loop that relates orders to backlog and delivery delay. In this loop, additional orders increase the order backlog and delivery delay. The longer delivery delay feeds back to reduce orders.

Figure 5

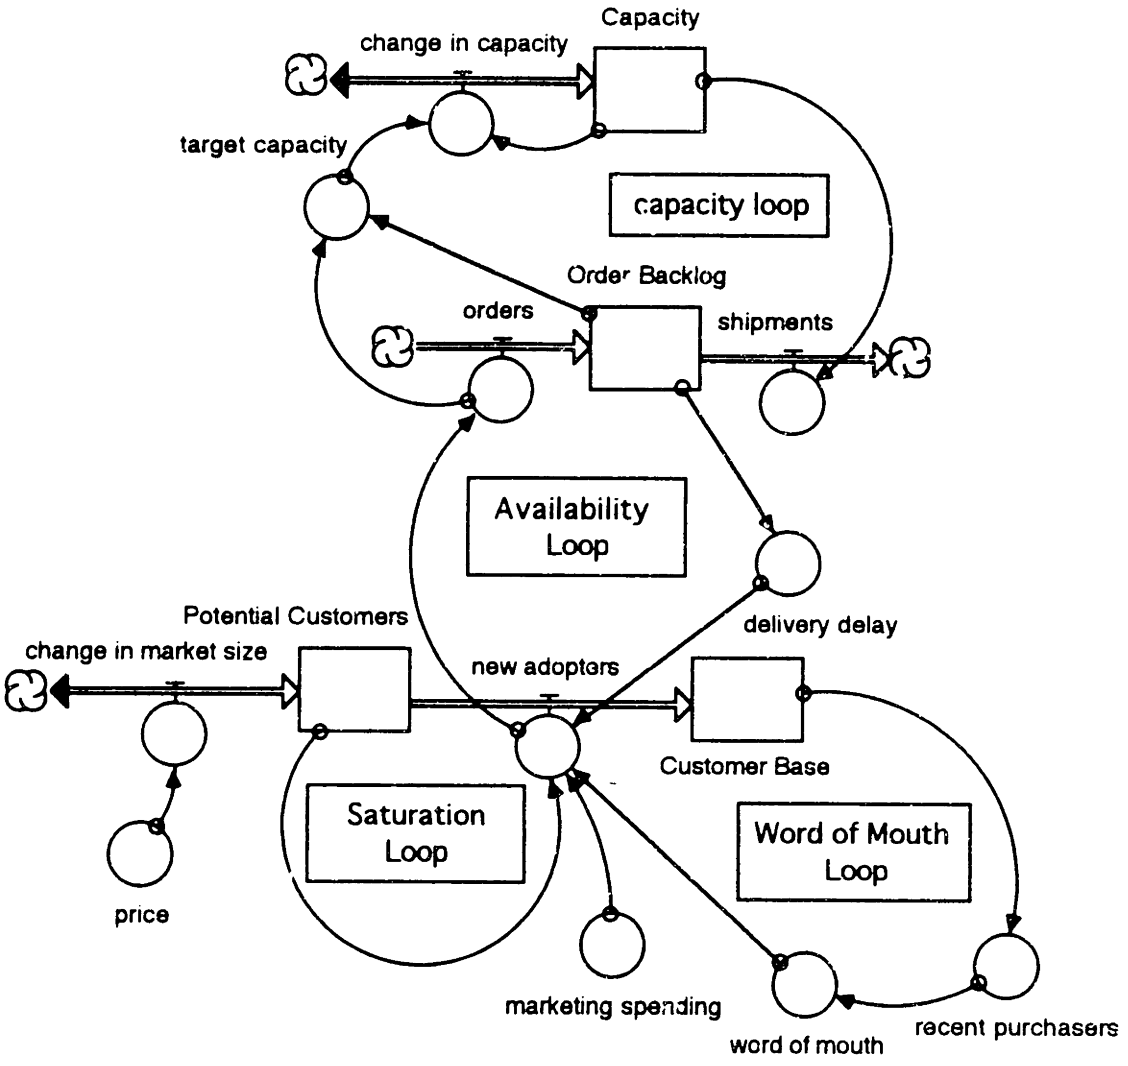

In the first example, the subject aggressively expanded capacity in response to order growth. This implies that the link between orders and target capacity was strong. Rapid capacity expansion reduced delivery delay during the growth phase but eventually resulted in excess capacity. In the second example, the subject expanded capacity slowly. This implies that the link between orders and target capacity was weak. Consequently, long delivery times reduced orders significantly. In the third example, the ballistic strategy, the link between orders and target capacity also was weak. The subject

preempted the market by determining target capacity based on the level of demand from previous games.

Benchmarks

Since potential profit in the scenarios depends on the strength of word of mouth and especially on replacement demand, the absolute profit levels cannot be used to compare performance across different scenarios.

To facilitate comparison across scenarios, two benchmarks were created that removed the effect of the treatment parameters on potential profit.

The first benchmark was the "perfect capacity adjustment" rule.

For each of the five scenarios, it was assumed that capacity was always equal to demand (the current backlog). There was never a shortage or excess of capacity which implied that delivery delay remained at its normal value of one month throughout the simulation.

In the "perfect capacity" benchmark, price is equal to constant markup over the sum of unit variable cost and unit fixed cost. For the constant markup rule, a 25% markup was approximately optimal for each of the five scenarios.

The "perfect capacity" benchmark controls for the size of the potential market. The subjects would be expected to perform better on some scenarios than on others regardless of the quality of their decisions. The reason for this is that the size of the potential market in a scenario depends on purchase frequency and the strength of word of mouth. The benchmark, which is the standard of comparison, should account for the effect of a larger market. Because there are no capacity constraints, the "perfect capacity" benchmark will always increase with the size of the potential market.

The second benchmark is the profit generated by simple, boundedly rational, decision rules for target capacity and price. A reasonable decision rule for target capacity must

include several factors. First, the rule must account for the size of the market. In a sensible rule, target capacity should increase with the size of the market.

Second, the decision rule must allow for market growth. It takes time for the firm to acquire or dispose of capacity. The delay in changing capacity means that the firm must anticipate changes in the market if it wants to match capacity and orders.

Third, the firm must consider its backlog relative to its capacity. Delivery times will be long if backlog is large relative to capacity. Long delivery times reduce market share and orders.

One reasonable decision rule for target capacity that includes these factors is shown below. This rule is similar in form to those used in Sterman (1989a,1989b) and Senge (1980):

1. $C^*(t) = s*D(t-1)*(1+g(t-1))^{a1}*(B(t)/C(t))^{a2}$
2. $g(t-1) = (D(t-1) - D(t-2))/D(t-2)$

where C* is target capacity, s is target market share, D is total market sales, g is the fractional growth rate of market demand, B is current order backlog and C is current production capacity.

The parameters can be interpreted as follows. The parameter a1 measures the firm's responsiveness to market growth. If a1 is large, growth expectations are important in the firm's target capacity decisions. The parameter a2 determines the firm's response when Backlog and Capacity are out of balance. A large value for a2 meant that the firm increases capacity significantly whenever backlog and capacity get out of balance.

The decision rule for price was based on a simple markup. Price is set as a 25% markup over the sum of unit variable cost and unit fixed cost. Price is responsive to neither the competitor's decisions nor to the supply/demand balance.

The decision rules for target capacity and price were incorporated into the simulation model and used to play all five scenarios. For ease of interpretation, the target market share, s, was set equal to .5, which was the initial value for the firm's market share in the simulation. The parameter values for a1 and a2 that maximized cumulative net income over all five scenarios were found through a grid search. Table 4 shows the cumulative net income generated by the decision rule for each scenario. The optimal parameters imply a conservative strategy. The value of a1 is 2.88 which means that target capacity grows more slowly than the market. This strategy causes market share to decline during the growth phase and then to rebound after demand begins to fall. The optimum value of a2 is .83 which means that target capacity increases by 8.3% when backlog is 10% greater than capacity.

Table 4

Benchmark Profit

Scenario	Perfect Capacity Rule	Naive Rule
1	$1,029,000,000	$634,119,000
2	$1,609,000,000	$1,039,580,000
3	$1,259,000,000	$922,088,000
4	$788,551,000	$229,138,000
5	$673,308,000	$397,515,000

The benchmark decision rules were unsophisticated and there is no claim that they were close to being globally optimal. The decision rules used locally available information and represent simple rules of thumb. No sophisticated information processing or forecasting schemes were considered. The rules did not even respond to

basic market forces of supply and demand or competitive price, much less engage in forward-looking strategic behavior.

Performance

Previous studies have shown that performance often is far from optimal in dynamic decision tasks and this study is no exception. Overall, the subjects performed very poorly relative to both benchmarks. The average ratio of the subjects' cumulative profit to the two baselines is shown in table 5. For all five trials, the subject's profit is 21% of the "perfect capacity" benchmark and 26.5% of the simple decision rule benchmark. Another performance measure is the simple decision rule's win/loss record against the subjects. Over all five trials, the model had higher cumulative net income with the naive rule in 492 of 565 trials for a win percentage of 87%.

	Table 5
Relative Performance Perfect Capacity Rule	.215
Relative Performance Naive Rule	.265

Above, it was hypothesized that performance should vary with the strength of the positive word of mouth feedback loop and the negative saturation loop. Strengthening either the positive or the negative loop is expected to make performance worse. This hypothesis was tested with the following general linear model:

3. MODRAT = CONSTANT + TRIAL + SUBJECT + LWOM + LRENT + LWOM*LRENT + TRIAL*LRENT + TRIAL*LWOM + v1

4. POTRAT = CONSTANT + TRIAL + SUBJECT + LWOM + LRENT + LWOM*LRENT + TRIAL*LRENT + TRIAL*LWOM + v2

where v1 and v2 are normally distributed random variables with zero means.

The variables MODRAT and POTRAT measured the subjects' performance relative to the two benchmark. MODRAT is a subject's cumulative net income for a particular scenario divided by the cumulative net income generated by the simple decision rules for that scenario. POTRAT is a subject's cumulative net income for a particular scenario divided by the cumulative net income generated by the "perfect capacity" decision rules.

The variables LRENT and LWOM are the natural logarithms of the parameters RENT and WOM respectively. A large value of LWOM implies a strong positive feedback loop connecting word of mouth communication and sales. A large value of LRENT means that consumers repurchase the product frequently and that the negative market saturation loop is weak. It was expected that both MODRAT and POTRAT would increase with LRENT and would decrease with LWOM.

The remaining variables are defined as follows. SUBJECT represents the individual subject who played the game. This variable controls for systematic differences between subjects' performance. TRIAL represents the trial number and it is expected that cumulative net income should increase over subsequent trials.

The results of the analysis of covariance are shown in tables 6 and 7. The trial effect is positive, statistically significant and shows that average performance improved during each trial.

			Table 6			
Dependent Variable: POTRAT						
	1	2	3	4	5	Significance
Trial	-.198	-.058	.004	.111	.141	.000
LWOM	-.121					.000
LRENT	.348					.000
Trial* LWOM	.117	-.047	-.029	-.026	.041	.561
Trial* LRENT	.018	.069	.009	-.003	-.067	.732
Subject						.001

			Table 7			
Dependent Variable: MODRAT						
	1	2	3	4	5	Significance
Trial	-.343	-.126	-.054	.214	.309	.000
LWOM	-.255					.001
LRENT	.611					.000
Trial* LWOM	.178	-.144	-.191	.045	.112	.517
Trial* LRENT	.214	.244	.036	-.196	-.298	.344
Subject						.050

The coefficients for LRENT and LWOM have the predicted signs and are statistically significant in both equations. Changes in both LRENT and LWOM are associated with large changes in the performance ratios.

The interaction effects between TRIAL and LRENT and TRIAL and LWOM are statistically insignificant for both benchmarks. The data do not support the hypothesis that the impact of feedback loop strength changed over the five trials.

To test the robustness of the results presented above, equation 3 was estimated with an alternative benchmark. The five scenarios each had a different profit potential because the market parameters generated different levels of sales. To control for the difference in potential profit, total market demand, as measured by units sold, was accumulated for each scenario. Cumulative net income was divided by cumulative unit sales for the scenario to create a profitability index (CSRAT). Dividing by cumulative sales controls for differences in market potential across the scenarios.

The results of estimating equation 3 with CSRAT as the dependent variable are shown in table 8. The results confirm the previous result that greater experience with the game does not attenuate the negative effect of feedback loop intensity on performance.

			Table 8			
Dependent Variable: CSRAT						
	1	2	3	4	5	Significance
Trial	-1.476	-.440	-.007	.819	1.104	.000
LWOM	-.894					.000
LRENT	2.397					.000
Trial* LWOM	.931	-.433	-.298	-.157	.043	.557
Trial* LRENT	.399	.664	.088	-.208	.943	.521
Subject						.003

Learning

Average performance improved significantly with experience but was still poor relative to both baselines. Even in the fifth trial, the subjects' average performance was just 57% of the naive decision rule's. The win/loss percentage also improved somewhat with practice. The naive rule won 76% of the games in the fourth trial and 82% in the fifth trial as opposed to 96% in the first trial.

The results demonstrate that the subjects did not learn to cope with the growth and saturation feedback loops. If the subjects had developed strategies more appropriate to an environment with strong feedback loops, the estimated coefficients on LRENT and LWOM would have become smaller over the five trials. Smaller values of LRENT and LWOM would have meant that the feedback loops had less effect on performance as the subjects gained experience with the game.

Other factors that improved performance over successive trials were accounted for by the direct effect of the TRIAL variable. For example, performance almost certainly improved because the subjects learned about the experimental set-up, the general size of the market and the approximate time the market would peak. In addition, performance could have improved as the subjects gained experience with the experimental set-up. These improvements were accounted for in the statistically significant direct effect of TRIAL.

The results of this study are consistent with and extend the misperception hypothesis. As expected, a stronger word of mouth loop or saturation loop is associated with large and statistically significant changes in performance. Somewhat surprisingly, the deleterious effect of feedback loop intensity does not diminish at all with experience even though the conditions for learning in the simulation are excellent. The subjects made 200 sets of decisions and received instant outcome feedback that was free of noise. There was no time pressure and the subjects had strong incentives for performance. Yet, in spite of these favorable conditions, the subjects did not learn to cope with environments involving strong feedback processes.

Appendix 1
Game Directions

MARKET STRATEGY: A Microcomputer Simulation of a Growth Market.

I. Introduction.

Market Strategy is an interactive computer game that simulates the dynamics of a rapidly growing market. The player assumes the roie of a manager who makes key decisions for a firm operating in a growth market. They player makes decisions for 40 quarters per game. Each player will play five different games that represent somewhat different market situations.

The player's objective is to maximize cumulative net income. Cumulative net income is defined as the sum of the net incomes earned over all forty quarters. Net income is defined in figure one. The various components of cost are discussed in more detail below.

II. Basic Structure of the Game.

The game is based on a simulated market environment. The player's firm sells a single durable product in a market with one competitor. The product was introduced recently and the firm has only two quarters of sales experience. The competitor introduced a similar product at about the same time.

The structure of the simulation is divided into three sectors, the market, the firm, and the competition. Each sector is described below.

A. The Market

Orders for the firm's product are determined by two components, total market orders and order share. Total market orders are the number of orders received by both the firm and the competitor. Order share is the fraction of total orders captured by the firm. Total market orders multiplied by order share equals the firm's order rate.

Total market orders are determined by several factors. It is assumed that there is a fixed population of households that would be interested in buying the product. In the simulation, this population is constant.

The total number of households that will eventually buy the product is determined by multiplying the population by the fraction willing to purchase. At current prices, only a fraction of the population of households would be willing to buy the product. The fraction that is willing to buy is determined by the minimum of the firm's price and the competitors price. For example, when the lowest price is $80, 10% of the population might be willing to buy. If the price were reduced to $40, perhaps 30% would be willing to buy.

Total market orders are reduced if many households have bought the product recently. Typically, in markets for durable goods, there is a significant delay between consumer purchases. For example, once someone buys a personal computer he is unlikely to buy a new one for several years. This process is replicated in the simulated market. Each consumer is assumed to buy one unit of the product and then wait several quarters before buying again.

Total market orders are influenced by two additional factors. First, word-of-mouth communication increases purchases. It is assumed that recent buyers help convince non-buyers to acquire the product. In short, recent sales lead to more sales. Further, it is assumed that the effect of recent sales decays over time. This represents the assimilation process whereby households eventually take their purchases for granted.

Second, marketing spending by both the firm and the competitor increases orders. Marketing spending includes total expenditures for advertising and promotion. It motivates consumers to purchase the product by providing them with information about product attributes. However, there is a delay before marketing spending has its maximum effect. This delay represents the time required to plan and implement marketing activities as well as the time required for advertising to "sink in."

In the simulation, the firm's share of total orders is determined by the attractiveness of the firm's product relative to the competition. Attractiveness is related to marketing spending, price, and delivery time. As would be expected, higher marketing spending and lower price both increase attractiveness.

Order share is also affected by delivery time. The firm's delivery time will be long if its order backlog is large relative to its production capacity. A long delivery time reduces market share.

Variables such as product quality, functionality, and price are assumed constant. These assumptions are made in order to simplify the structure of the game. The firm can influence sales only by controlling price and target capacity.

B. The Firm

1. Orders, Backlog, Shipments.

Figure two shows the firm's backlog system. Orders for the firm's product accumulate in an order backlog until they can be shipped. The firm will ship its entire order backlog, less cancellations, if it has adequate production capacity. Otherwise, it will ship an amount equal to its production capacity. For example, if the firm has 20,000 units of capacity and a backlog minus cancellations of 10,000 the firm will ship 10,000 units. If the firm has a backlog of 30,000 units, it will ship 20,000 units.

Each quarter, the firm loses a fraction of its orders due to cancellations. The cancellations take place because people get tired of waiting for their merchandise.

2. Capacity

Figure three shows the structure of the firm's capacity acquisition process. Target Capacity is the firm's desired level of productive capacity measured in units per quarter. At the beginning of a quarter, the firm has a stock of plant and equipment available for production. The stock of plant and equipment translates into the maximum amount that the firm can produce during a quarter. The firm can never produce more that its capacity.

It takes time for the firm to adjust capacity to its target level. In the game, it takes an average of four quarters to acquire a unit of capacity. However, the four quarters is an average, some capacity can be added sooner while other expansion projects take longer.

Figure four shows how capacity responds to a one-time increase in target capacity. At quarter 5, target capacity increases from 1,000 units to 2,000 units per quarter. Capacity immediately starts to grow as the firm does the easiest expansion projects first. Capacity gradually approaches target capacity as the firm completes longer, more difficult projects.

3. Control Variables

The player can control orders and shipments by manipulating price and target capacity. Each quarter, the player will enter decisions these two variables. The decisions will be used to calculate subsequent values of the model variables.

The amount of marketing spending is not determined by the player. Marketing Spending is set at $50,000 for the first two quarters. After the second quarter, marketing spending is always set at five percent of the last quarter's revenue.

4. Revenue / Costs

Revenue is determined by sales and the average price per order. In the model, price changes from quarter to quarter. The average order price is a weighted average of past prices.

The components of the firm's costs are shown in figure one. "Variable Cost" is determined by orders and unit variable cost. Unit variable cost starts at $30 per unit and falls as the firm's cumulative sales increase. The reduction in unit variable cost caused by higher cumulative sales simulates the behavior of a learning curve.

"Fixed Cost" is determined by capacity and unit fixed cost. It is assumed that each unit of capacity generates debt service and overhead cost each quarter. Unit fixed cost is constant at $30 per unit of capacity.

"Investment Cost" is what the firm must spend when it increases or decreases capacity. These costs include installation costs and the cost of scrapping capacity. Capacity adjustment cost is determined by the change in capacity during a quarter and the unit adjustment costs. The unit adjustment cost is constant at $45.

C. Competition

As mentioned above, the firm faces a single competitor. The competitor's decisons for target capacity and price are made by the computer. The player is provided with information about the competitor's price, marketing spending, and capacity. However, the player is not provided with information about the rules that detemine these decisions.

The competitor faces exactly the same physical and information limits as the firm. Also, the competitor's costs are determined in the same way as the firm's.

III. The Game Board.

To begin the game, insert the game disk into the MacIntosh disk drive. After the MacIntosh starts up, a file with a "Penguin" icon will be displayed. Open this file to start the game.

After the file in opened, you will be asked to enter your name and the trial number you are beginning. If you have started the game for the first time, enter "1" for the trial number. If you have finished trial number one and are starting trial number two, enter "2" for the trial number. After you have entered your name and the trial number, the game board will be displayed.

The game board is laid out like a spreadsheet. The columns of the spreadsheet are quarters while the rows correspond to various important quantities. Each of the variables displayed on the spreadsheet is defined in figure one.

Some of the variables that can be displayed on the game board are not immediately visible because of limited screen size. To access these variables, use the scroll bar on the left side of the screen to move down the rows of the spreadsheet.

The bracket symbol (< >) indicates a decision variable. The player must decide on target capacity and price. Before the player makes his decision, there will be brackets in the boxes for those variables in quarter three. The players decisions in quarter three, along with prior decisions, will determine revenue, cost and net income for quarter three. These decisions will also determine the stock of capacity available at the beginning of the fourth quarter.

IV. Entering a Decision.

In order to enter your decision for price, click the mouse on the round button that is next to the label "Price". After you have clicked on "Price", the button will turn black. You can now type in your decision. The number you typed should appear in the box in the upper left hand corner of the screen. To enter your decision into the computer, simply press the return key. Your decision will replace the bracket symbol and appear in the "Price" box for quarter three

After you press return, the round button for "Target Capacity" will turn black. You are now ready to type in your decision for "Target Capacity". After you have typed in

your decision press the return key. Your decision will appear in the box for quarter three "Target Capacity".

If you wish to change your decision for "Price", press the return key to move back to the "Price" row. Pressing the return key will move you back and forth between "Price" and "Target Capacity".

Your decisions for quarter three can now be used to calculate the results for quarter three and the value for capacity at the beginning of quarter four. To calculate the results for the quarter, click on the button labeled "Calculate Quarter". Alternatively, results can be calculated by pressing the "Tab" key.

If you mistyped either of your decisions, you can reenter your decisions by clicking on the "Backspace" button. Clicking on "Backspace" will roll the game back to beginning of the previous quarter and allow you to correctly enter your decisions.

In summary, follow these steps to enter a decision.

1. Click on the round button in the row for "Price".

2. Type in your decision for "Price". The number should appear in the box in the upper left hand corner of the screen.

3. Press the return key. Your decision for "Price" should appear in the box for "Price" in the appropriate quarter.

4. Type in your decision for "Target Capacity". The number should appear in the box in the upper left hand corner of the screen.

5. Press the return key. Your decision for "Target Capacity" should appear in the box for "Target Capacity" for the appropriate quarter.

6. Check to make sure the correct decisions are entered.

7. If there is an error, press the return key to move to the appropriate box and then correct the error.

8. If there are no errors, press the "Tab" key or click on the button labeled "Calculate Quarter".

V. Moving Between Quarters in the Spreadsheet.

To move between quarters in the spreadsheet, you can use the scroll bar on the top of the screen. The number next to the scroll bar is the quarter displayed on the left-hand side of the spreadsheet. Drag the scroll bar back and forth in order to display different quarters. You can move to any of the 40 quarters. If you want to return to the current quarter, click on the box labeled "Go to Current".

VI. Plotting.

Any variable shown on the spreadsheet can be plotted over time. In order to plot a variable, click on the square button next to the variable name. The button will turn black after you click on it. To display the graph, drag open the menu labeled "Options" and click on the heading labeled "Show Graph". The graph will be displayed. To return to the spreadsheet, drag open the menu labeled, "Options" and click on the heading labeled "Hide Graph". To remove a variable from the graph, click on the square graph button. A maximum of four variables can be shown on the graph.

For quantities that are always greater than zero, the minimum value of the graph's scale is always zero. The maximum value will be the largest value of the variable through that quarter. For quantities that go below zero, an appropriate minimum value will be automatically selected. Negative numbers are enclosed in parenthesis.

The scales for the variables are shown at the bottom of the screen in the form (minimum value/maximum value/scale interval). For example, "Sales(0/20000/4000)" means that the minimum point on the sales graph is 0, the maximum point is 20,000, and the interval between "tick marks" is 4,000. For very large numbers, the scales values are in exponential notation. The plotted variables can be constrained to be on the same scale by clicking on the "Plot Same Scale" button at the top of the spread sheet.

VII. Starting a New Game.

The game consists of 40 quarters of decisions. After you have completed 40 quarters, the computer will automatically save your results and give you instructions to

start a new trial. At the end of five games, the computer will tell you that the session is finished and give you instructions to quit.

VIII. Quitting the Game.

You may quit at the end of any trial and begin a new trial at another time. To exit the game program, drag open the menu labeled "File" and select the heading labeled "Quit".

However, **please do not quit in the middle of a trial**. The game software will assume that you have completed the trial and move you on to the next one.

To restart the game, click on the game icon and enter the trial number that you will start with. For example, if you quit after finishing two trials, enter trial number 3 when starting the game is restarted.

Chapter Six
Discussion of Results

The first section of this chapter analyzes the results presented in chapter three. The second section considers the learning that took place over the course of the simulations. The third section discusses the results as an example of the misperception of feedback.

I. Analysis of the Experimental Results

The experiments described in chapter three yielded four primary results. First, even after considerable experience with the task (200 decision rounds), the subjects performed very poorly in comparison to both the naive and the perfect foresight baselines. Second, the subjects' performance deteriorated significantly when either the self-reinforcing word of mouth feedback loop or the balancing repurchase feedback loop was strong. Third, with respect to learning, the subjects' average performance level improved significantly with practice; there seems to be a strong general learning effect although not strong enough to bring the subjects' performance up to the benchmarks. Fourth, the deleterious effect of feedback loop intensity on performance was not mitigated by experience.

Specification of the Decision Rules

Characterizing the subjects' decision making is the first step in the process of analysis. Previous studies of dynamic environments have used linear decision rules to characterize the outcome of subjects' decision making processes (Sterman 1989a, 1989b, Mcann, Miller & Moskowitz 1975, Remus 1981, Sterman & Diehl 1992). The form of the estimated rules is usually very simple and there is no claim that the subject actually applied the rule in making decisions. Rather, a well specified rule describes the

outcomes generated by the more complex cognitive processes that convert information inputs into actual decisions (Sterman 1989b).

The specification of the decision rules is, in part, suggested by the structure of the task. In the simulated market, the basic pattern of demand was rapid growth followed by a peak and then decline to a stable equilibrium. Although the growth rate, the timing of the peak, and equilibrium demand all differed between scenarios, the fundamental pattern of growth and decline was the same in all five scenarios.

For the capacity decision, the subjects' decision problem can be broken into two stages. The first stage was managing capacity through the growth and decline phase. In this phase, the DMs needed to balance the risk of overexpansion against the opportunity cost and lost profit if capacity from insufficient capacity during the growth phase.

The second problem was managing capacity during the equilibrium phase of the market. In the simulations, the market usually peaked between quarters twelve and twenty. Following the peak, the market declined to a stable level for the remainder of the simulation. The DMs task was to determine the combination of capacity and price that earned the largest profit. Since the market was not growing (or growing only slowly), the task was equivalent to choosing a target market share.

A rule that incorporates both aspects of the decision problem will be developed in the following sections. The functional form of the rule is:

1. $C^*(t)=S^*[De^{(1-\beta 1)}D(t-1)^{\beta 1}](1+g(t-1))^{\beta 2}(B(t)/C(t))^{\beta 3}$

2. $g(t) = (D(t-1)-D(t-2))/D(t-2)$

where,

C* - Target Capacity

De - Prior Expectation of Average Market Demand

S - Target Market Share

D - Total Market Demand

g - Last quarter's growth rate of market demand

B - Order Backlog

C - Production Capacity

ß1,ß2,ß3 - Parameters

The first three terms in equation 1 imply that Target Capacity was determined by setting a target share S of the weighted average of expected equilibrium demand De and last quarter's market demand D(t-1). The weighting parameter, ß1, measures the willingness of the subjects to increase capacity when market demand is above its expected long-run level. If ß1 equals one, the DM based target capacity only on current demand. If ß1 equals zero, the DM ignores market demand, never increasing capacity above the expected long-run equilibrium level.

The fourth term in equation one, 1+g(t), captures the response to demand growth. During the growth phase, because there is a four quarter lag in capacity acquisition, target capacity must increase in advance of actual market demand in order to maintain a constant market share. For example, if ß2 equals zero, the DM ignores actual market growth in determining target capacity. If ß2 equals 2, target capacity includes an allowance for two quarters of expected market growth.

The fifth term, (B(t)/C(t)), represents the response to delivery delay. The ratio of the order backlog to production capacity is the firm's average delivery delay, a measure of the demand/supply balance. For example, if backlog is fifty and capacity is ten, it takes an average of five time periods before a new order can be shipped. A larger value of ß3 , implies that the DM was more responsive to long delivery delays.

The proposed decision rule for price is the product of unit variable cost (UVC) and a target markup (TM). This type of simple cost-markup rule is widely used in real firms as well as models in both system dynamics and economics (Lilien & Kotler 1992).

The target markup depends on a measure of the firm's supply-demand balance and the firm's policy for passing cost reductions on to the consumer. It would have been reasonable for target markup to increase when demand was large relative to supply. One measure of the supply-demand balance is the ratio of order backlog to capacity. If backlog is large relative to capacity, delivery delay is long and the DM might decide to raise price. The DM might cut price if backlog is small relative to capacity.

The DM also decides whether to pass cost reductions on to the consumer. In the simulations, unit variable cost fell as cumulative volume increases. The DM faced a choice of whether to reduce price or to maintain price and keep the cost reduction as profit.

A price equation that incorporates these factors is:

3. $P(t) = UVC(t)*M(t)$

4. $M(t) = Mo*(B(t)/C(t))^{\beta 4}(UVC(t)/UVC(0)^{\beta 5})$

where UVC(0) is the initial value of UVC.

Substituting equation 3 into equation 2 yields:

5. $P(t) = K*(UVC)^{(1+\beta 5)}(B(t)/C(t))^{\beta 4}$

where

$K = Mo/(UVC(0))$

Mo - Baseline Markup

M - Markup

UVC - Unit Variable Cost

ß4,ß5 - Parameters

The parameters of the price rule can be interpreted as elasticities. The parameter determined how much of the cost reduction was passed through in a lower price. If ß5= -1, the DM used a skimming strategy that takes all the cost reduction as increased profit margin. If ß5=-.5, a ten percent reduction in unit cost causes a five percent reduction in price. As increases (becomes less negative), a larger share of the cost reduction is passed through as a lower price. Larger values correspond to an aggressive learning curve strategy in which the DM attempts to build market share through lower price

The parameter ß4 determined the responsiveness of price to the supply-demand balance. A large value of ß4 means that the DM aggressively changed price to balance capacity and backlog. A small value corresponds to a DM who balanced supply and demand with delivery time instead of price.

The price equation does not include a direct response to competitor price. Competitor price was included in some preliminary versions of equation four but was dropped due to extreme collinearity between competitor price and unit variable cost. The competitor's price is a constant markup over unit variable cost and the initial value for the competitor's unit variable cost is equal to the subject's. Although the learning curve allows the subject's cost to diverge from the competitor's, they were usually very close which creates a strong correlation between competitor price and the subject's unit variable cost.

The price equation indirectly accounts for the competitor's actions. If the competitor cuts price and reduces demand for the subject, the ratio of backlog to capacity could fall. In this case, equation four would dictate a price cut. If the competitor cuts price, and backlog does not fall relative to capacity because orders are still greater than capacity, equation four would not generate a price cut. This response is sensible for there would be no gain from cutting price if existing demand cannot be satisfied.

Estimation of the Decision Rules

In equation one, it was not possible to estimate the dynamic processes that determined expected equilibrium demand and target share. In reality, both expected equilibrium demand and target share were functions of the information cues generated in the simulation. For example, the subjects had some prior expectation of what equilibrium demand would be and then modified that expectation based on the market demand they actually experienced (Sterman 1988). Unfortunately, the data were not sufficient to estimate the parameters of the underlying process that determined expected equilibrium demand and target share. Instead, both variables were assumed constant and are aggregated in the constant term (a0).

To estimate equations one and four, an additive error term was assumed, which yielded the following equations,

6. $\log(C^*(t)) = ao + a1*\log(D(t-1)) + a2*\log(1+g(t)) + a3*\log(B(t)/C(t)) + u1$
7. $\log(P(t))=bo+ b1*\log((UVC) + b2*\log(B(t)/C(t)) + u2$

where u1 and u2 are normally distributed random variables with zero means.

Ordinary least squares estimates of equations five and six showed that there was significant serial correlation in the error terms. This result was not surprising for there are at least two factors that could produce autocorrelation. First, it is likely that the DMs used past decisions as an anchor in formulating current decisions. Excluding past decisions from the estimated equations will produce autocorrelation. Second, the decision rules are only approximate descriptions of the real decision process and often leave out relevant factors. The factors omitted from the decision rule cause serial

correlation if, as is likely, they are correlated over time (Mcann, Miller, Moskowitz 1975).

Consequently, the Cochrane-Orcutt procedure was used to correct for serial correlation. The Cochrane-Orcutt procedure assumes that a first order auto-regressive process generates the error term and estimates the auto correlation coefficient. The estimated auto correlation coefficient is used to make the appropriate corrections to the estimated parameters (Pindyck and Rubinfeld 1989).

Both decision rules were estimated for each of the 568 trials in the sample. Each trial included the 38 decisions made by the subject for target capacity and price.

Table one shows the results of the estimation. Generally, the models do a good job of explaining the subjects' capacity and price decisions. The median R2 is .87 for the target capacity rule and .94 for the price rule. Table one shows the percentages of each coefficient that are significant at the 5% level or better.

Table 1

Parameter	Mean	Std Dev.	% Significant at 5% or better	% Expected Sign
Capacity Rule				
a0	8.414	6.45	79%	93%
a1	.383	.433	71%	80%
a2	.036	.533	44%	45%
a3	.318	.715	57%	68%
Rho	.560	.328	79%	
R2	.872	.183		
Pricing Rule				
b0	3.125	2.553	97%	97%
b1	.259	.337	72%	85%
b2	.016	.067	38%	69%
Rho	.781	.215	97%	
R2	.947	.095		

As shown in table one, the majority of the coefficients have the expected signs. The exception is the coefficient for expected growth, a2, for which only 45% of the coefficients are greater than zero. One possible explanation for the negative values of a2 is that subjects adopted a counter-cyclical strategy. Over the course of the trials, the subjects could have learned that rapid growth is a harbinger of rapid decline. In this case, rapid growth could have provided a signal to reduce target capacity which implies a negative value for a2.

Obviously, the growth coefficient is most important in scenarios where the market grows rapidly. For example, in the extreme case, the growth coefficient would be irrelevant if there was no market growth. Because the growth coefficient, a2, is more important in the growth scenarios, a2 should be larger when growth is fast. This implies that a2 should be related to the parameter, WOM, that controls the market growth rate.

Assume that the subjects chose between a pro-cyclical strategy and a counter-cyclical strategy. In the pro-cyclical strategy, a2 would be positive and in the counter-cyclical strategy a2 would be negative. Because a2 is most important in rapid growth situations, it is expected that a larger value of WOM would result in a larger positive number for a2

with pro cyclical strategy and a larger negative number for a2 with the counter-cyclical strategy.

This expectation was tested by regressing LWOM and a set of control variables against a2. A categorical variable (CCYC) was set equal to 1 if $a2 > 0$ (the pro-cyclical strategy) and equal to 0 if $a2 < 0$ (the counter-cyclical strategy). As shown in table two, a larger value of LWOM is associated with the pro-cyclical strategy and a reduction in a2 (makes it more negative) with the counter-cyclical strategy. The relationship between a2 and LWOM is consistent with the hypothesis that the subjects made a choice between a counter-cyclical and a pro-cyclical strategy.

	Table 2	
Dependent Variable: a2		
Parameter	Coefficient	Significance
LWOM	-.299	.007
CCYC*LWOM	.325	.000
RENT	.036	.305
TRIAL 1	-.103	.026
TRIAL 2	-.045	
TRIAL 3	.012	
TRIAL 4	.114	
TRIAL 5	.022	
R2	.091	

Evolution of the Estimated Coefficients

The coefficients of the target capacity rule change significantly over the five trials. Table three shows the result of an analysis of variance that related decision rule coefficients to the trial, the subject, and the treatment parameters. As shown in table three, the coefficient a0 became significantly larger over the five trials, a1 became significantly smaller, a2 grew larger, and a3 changed significantly but did not evolve in a consistent direction over the five trials. Surprisingly, none of the coefficients in the price rule changed significantly over the five trials.

Table 3

Coefficient	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Significance
a0	6.601	8.178	8.32	9.29	9.602	.001
a1	.485	.392	.397	.331	.315	.000
a2	-.064	-.015	.053	.14	.086	.018
a3	.099	.317	.255	.487	.347	.000
b0	3.39	3.01	3.00	3.27	2.92	.502
b1	.26	.27	.24	.30	.23	.489
b2	.01	.01	.02	.02	.02	.381

Cluster Analysis

The changes in the coefficients for the target capacity rule suggested that the subjects were systematically altering their capacity rules over the five trials. In order to capture the evolution of the rules, cluster analysis was used to group the estimated decision rules into groups that represent distinct price and capacity strategies. In previous studies, cluster analysis has been used to classify firms' strategies into groups according to variables like marketing intensity and relative price (Harrigan 1985). In this instance, the classification is based on the standardized parameters of the subjects' decision rules.

The specific clustering technique employed was KMEANS cluster analysis which assigns observations (in this case, the trials) to clusters so as to maximize the variance between clusters (Hartigan 1975). KMEANS requires the analyst to choose the number of clusters to be estimated and there are no set criteria for making this choice. In this case, the criterion was to choose the smallest number of clusters that revealed

fundamental differences between the strategies. After some experimentation, it was determined that five clusters was the smallest number required to isolate meaningful strategies. Three of the five clusters accounted for 536 of the 565 trials (95%) so most of the subsequent analysis will concentrate on these groups. Table four shows the non-standardized value for the mean for each estimated parameter and the number of trials included in the three dominant clusters.

	Table	**4**	
	Cluster 1	Cluster 2	Cluster 3
Coefficient			
a0	13.604	3.13	8.822
a1	.052	.743	.309
a2	.063	-.168	.044
a3	.047	.607	-.085
b0	3.337	3.545	2.955
b1	.112	.292	.529
b2	.044	.024	-.062

The clear differences between the strategies captured by the clusters can be illustrated by running the simulation model with the mean parameters for clusters one and two and the treatment parameters for scenario one. Figures 1 and 2 show the behavior of target capacity and capacity for the two strategies. In cluster two (C2), target capacity tracks the market as demand increases slowly, overshoots and then declines. In this strategy, which will be called the tracking strategy, target capacity responds to the growth and decline of market demand although it lags behind. The cluster one (C1) strategy is much less responsive to current market demand than the C2 strategy. For C1, target capacity

quickly increases and then stabilizes. Target capacity is unresponsive to market demand and never goes very far above its equilibrium level. The C1 decision rule could be referred to as a "ballistic strategy" in that target capacity does not change in response to feedback from the market.

Figure One

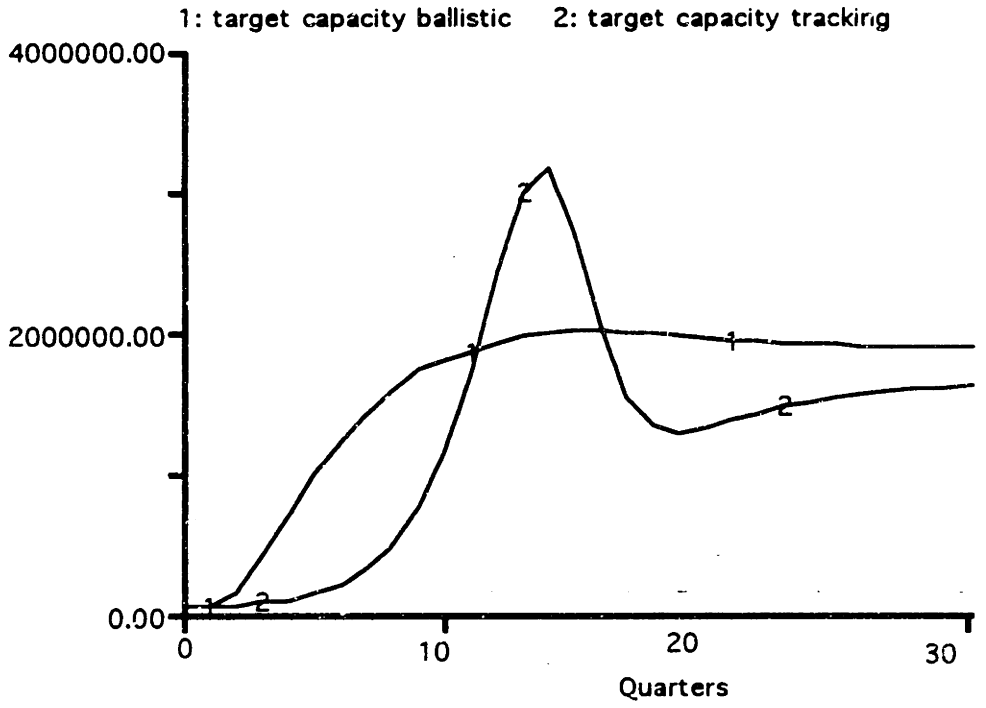

Figure 2

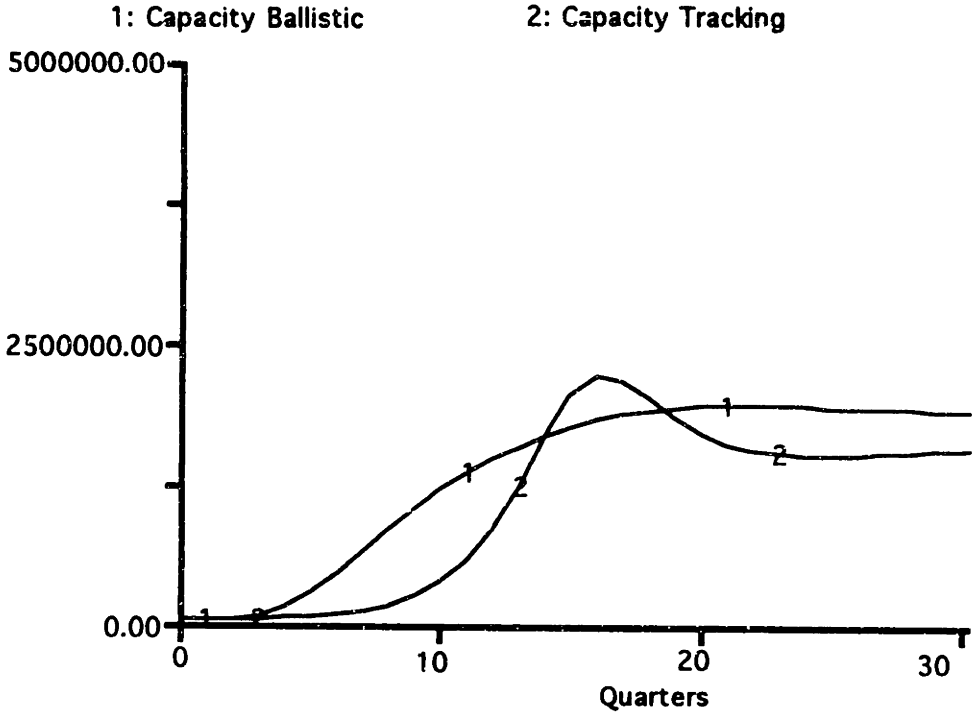

As shown in table five, the observations in the clusters are not distributed randomly across the five trials. The percentage of trials in cluster one (the ballistic strategy) increased from 26% in the first trial to 51% in the fifth trial. The number of trials in the other clusters showed a corresponding decrease. The cluster analysis confirms that the coefficients of the decision rule systematically evolved toward a strategy that used less outcome feedback from market demand.

		Table 5	
Trial	Cluster 1	Cluster 2	Cluster 3
1	26.42%	42.45%	31.13%
2	41.23%	40.35%	18.42%
3	43.81%	42.86%	13.33%
4	50.49%	33.01%	16.50%
5	50.93%	32.41%	16.67%

Of course, the cluster analysis does not imply that the subjects literally made discrete jumps from one cluster to another. Rather, over subsequent trials, the subjects gradually moved in the direction of the ballistic strategy. Dividing the trials into clusters is a convenient means for summarizing and analyzing the continuous evolution of the decision rules.

II. Learning in the Life-Cycle Simulation

It was expected that two types of learning would take place during the experiment. First, there is general learning about the experimental set-up, the basic pattern of market demand, and the behavior of the competitor. For the first trial, the subjects have little idea what to expect. For example, although the subjects had been exposed to product life cycle concepts, they had no idea where the market would peak within the forty quarters or whether it would continue to grow. After the first two trials, the subjects could reasonably expect that they were dealing with a situation of growth and decline.

The second type of learning is the subjects' acquisition of more effective capacity and price strategies. People would be expected to learn gradually what strategies are most successful in the life-cycle setting. In the best of circumstances, the subjects would learn to apply robust decision rules that use outcome feedback to adapt to a wide variety of specific demand patterns.

The relative magnitude of the two types of learning during the five trials was tested by estimating the following relationship

8. POTRAT= CONSTANT + TRIAL + STRATEGY + LWOM + LRENT + TRIAL*LWOM + TRIAL*LRENT + STRATEGY*LWOM + STRATEGY* LRENT

The first five terms of equation eight account for mean performance and the main effects of general learning, the choice of strategy, and the treatment parameters. The trial variable should account for all the sources of general learning other than the choice of strategy. The categorical strategy variable denotes which strategy, the ballistic or either of the two tracking, the subjects used in each trial (assigned according to the cluster analysis). The first tracking strategy (cluster 2) is more responsive to changes in market demand and the supply/demand balance than the second tracking strategy (cluster 3). The strategy variable should estimate the average effect on performance of choosing one of the three strategies. It was expected that the trial effect would be significant and would show that performance improved with experience. The strategy variable was expected to be significant but it was unclear which strategy produced the best performance.

The next four terms account for interaction effects. The interaction between trial and the treatment parameters indicates whether general learning changes the effect of feedback loop intensity. Based on the results from chapter three, this effect was expected to be insignificant. The interaction between the strategy and the treatment parameters indicates

whether the choice of strategy mitigates the effect of feedback loop intensity. It was uncertain whether this interaction would be significant. The final interaction is between trial and strategy. There was no strong expectation about whether the performance of any of the strategies would improve with experience.

The results of estimating equation 8 are shown in table six . As expected, the trial effect and the treatment effects were large and significant. For the strategy variable, the ballistic strategy (C1) performed significantly better than either of the tracking strategies. However, the estimated general learning effect was much larger than the effect of strategy choice. Over the five trials, performance relative to the perfect foresight baseline improved by an average of 28.2%, while the choice of the ballistic strategy instead of the tracking strategy produced an improvement of 15.5% of baseline profitability. The general learning effect was 82% greater than the effect of strategy choice.

Table 6

	1	2	3	4	5	Significance
Constant	.207					
Trial	-.18	-.073	.043	.108	.102	.000
Strategy	.093	-.062	-.031			.001
LWOM	-.102					.002
LRENT	.323					.000
Trial* LWOM	.034	.011	-.010	.019	.054	.916
Trial* LRENT	-.008	-.013	.022	.057	-.058	.821
Cluster* WOM	.008 (Ballistic)	-.132 (Tracking 1)	.124 (Tracking 2)			.009
Cluster* RENT	-.02	.061	-.041			.373

The only significant interaction effect was between the strategy and the word of mouth parameter. The negative impact of the word of mouth parameter was significantly greater with the adaptive strategy than with the ballistic strategy. This indicates that the ballistic strategy eliminated some, but not most, of the effect of the word of mouth feedback loop.

None of the three strategies mitigated the very large negative effect of the market saturation parameter.

The results of estimating equation eight have implications for the nature of learning in the experiment. First, the results show that the ballistic strategy outperformed the adaptive strategy but that both are very poor relative to the naive baseline. Even in the fifth trial, the average net income generated by those who used the ballistic strategy is only two-thirds of the simple baseline. Also, adopting the ballistic strategy does nothing to reduce the negative impact of the market saturation loop and reduces, but by no means eliminates, the effect of the word-of mouth loop.

Explanation for the Performance of the Ballistic Strategy

To explain the improved performance generated by the ballistic (C1) strategy, a series of simulations was run to compare the performance of typical C1 (ballistic) and C2 (tracking) decision rules. The decision rules selected for the simulations used the mean parameters for each cluster. The decision rule parameters varied significantly within the clusters but using the mean parameters captures the fundamental differences between the two strategies. The strategies were compared using scenario four which was the most difficult of the five. In scenario four, market demand grows quickly, peaks in quarter 12 and quickly declines to a relatively low equilibrium level.

The simulations reveal two reasons why performance was better with the ballistic strategy. First, the ballistic strategy quickly increased capacity to the expected equilibrium level in advance of market demand and captured additional profit. This can be seen in figure three in which net income with the ballistic initially grows much more quickly than with the tracking strategy.

Figure 3

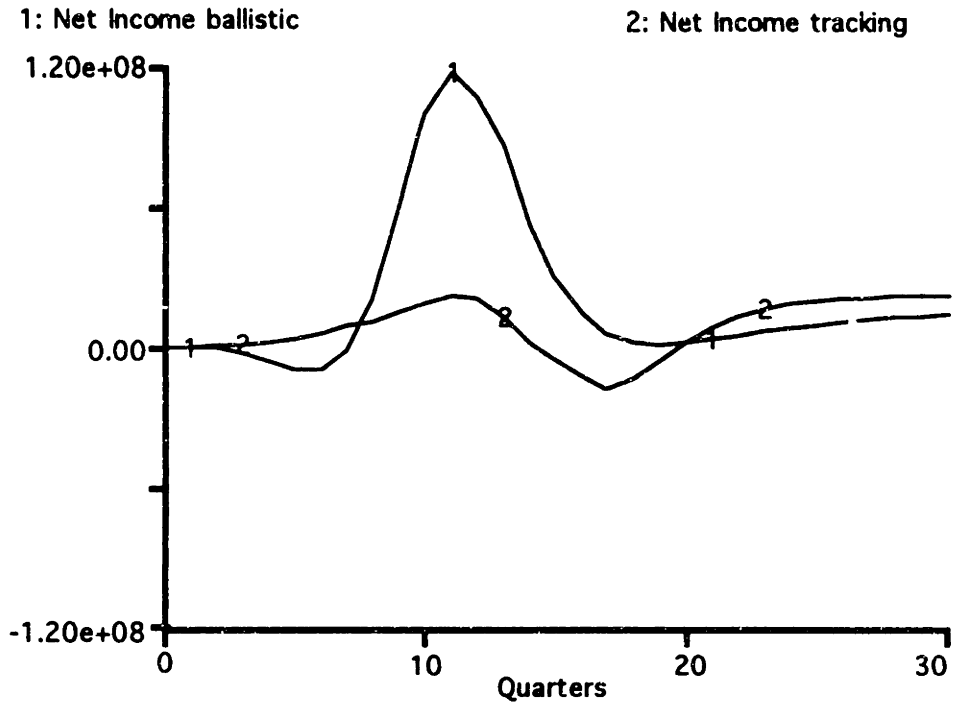

Second, the ballistic strategy did not attempt to build capacity to meet peak demand. Capacity rapidly increased to the expected equilibrium level of sales and then stopped, which effectively chopped off the peak of market demand. After market demand started to fall, there was much less excess capacity in the ballistic strategy because capacity never reached as high a level. Also, because capacity didn't change as much in the ballistic strategy, either up or down, adjustment costs are significantly lower. Figure 3 shows that the ballistic strategy incurred much smaller losses during the market decline.

Overall, the ballistic strategy produced superior results because it better exploited the market's growth phase and avoided excess capacity. These dynamics are robust in the sense that selecting alternative sets of parameters from the clusters would not change the comparative results significantly.

However, despite the average improvement, the ballistic strategy still performed relatively poorly in the difficult scenarios. The simulations revealed two reasons for this result. First, even though the ballistic strategy based target capacity mostly on prior demand expectations, it still responded slightly to market demand and backlog. This caused some excess capacity and adjustment costs. Second, and most important, in scenario four, which was the most difficult, most of the potential sales were concentrated during the period of rapid market growth. The ballistic strategy did not expand capacity to satisfy peak demand and ceded these sales to the competitor. The sales and profit foregone were larger for scenarios four and five because they have sharper peaks. Relative profit was lower in these scenarios because the ballistic strategy gave up more of the potential profit to the competitor than in scenarios with flatter peaks. The ballistic strategy accepted lower sales in exchange for less excess capacity. On the average, this exchange increased profit but did not eliminate the difference in relative performance between the easier and more difficult scenarios.

Outcome Feedback in the Simulated Market

In the first few game trials, the subjects used outcome feedback about the level of market demand to determine target capacity. Prior to the first trial, the subjects received information about the system's structure but had little knowledge about the pattern or level of demand. Based on their poor performance in the initial trials, the subjects altered their strategies to use less, not more, outcome feedback. The subjects used their experience from previous trials to anticipate the approximate size of the replacement market and quickly expanded capacity to that level. The increase in target capacity was independent of the level or growth rate of market demand actually experienced in the specific scenario.

The evolution of the subjects' strategies highlights the distinction between "within-trial outcome feedback" and "between-trial outcome feedback." In the later trials, relatively little within-trial outcome feedback was used. Target capacity was unresponsive to either the level or the change in market demand and to the size of the order backlog. Instead, target capacity was based primarily on between-trial outcome feedback. The subjects learned about the pattern of demand from previous trials and paid little attention to outcome feedback in the particular trial.

The ballistic strategy fell far short of the performance levels that are feasible with the appropriate use of within-trial outcome feedback. To analyze the appropriate use of outcome feedback, a series of simulation experiments were conducted with the "best" parameters for equations 5 and 6. The "best" parameters (which are shown in table 8) generated the highest cumulative net income over all five scenarios and were determined by a grid search, an examination of the subjects' most successful decision rules, and the author's experience with the simulation model. There is no claim that the "best" parameters are optimal, but, as shown in table seven, they generate cumulative net income that is three times greater than the average subject and more than twice as large as the average of subjects who used the ballistic strategy.

Table 7

Strategy	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Total
Average Subject	$295	$593	$605	$-120	$48	$1,418
Subject Using Ballistic Strategy	$368 millions	$652	$809	$45	$106	$1,908
"Best" Strategy	$882	$1,328	$986	$684	$511	$4,391

The superior performance of the "best" rule can be explained by the appropriate use of outcome feedback. Table eight compares the performance generated by the "best"

parameters and the parameter of the average ballistic rule. The "best" rule places significantly more weight on the outcome feedback variables than the average ballistic rule. The weights on the outcome feedback variables, a1, a2, a3, and b2, are an average of 622% higher in the "best" rule than in the average ballistic rule.

Table 8

Strategy	a0	a1	a2	a3	b0	b1	b2
Average Ballistic Strategy	13.6	.052	.063	.047	3.33	.112	.044
"Best" Strategy	11.5	.166	1.65	.5	3.67	.098	.151

Stronger outcome feedback explains the performance improvement produced by the "best" strategy. First, stronger feedback from backlog and market growth causes capacity to track demand more closely. During the growth phase, closer tracking reduces delivery delay and increases market share. During the decline phase, closer tracking reduces the cost of excess capacity. It was expected that adding outcome feedback from growth and backlog would help correct gaps between capacity and orders. For example, if capacity was large relative to demand there would be pressure to reduce target capacity. Simulations of both strategies, for all five scenarios, revealed that the "best" strategy drastically reduced the average gap between orders and capacity as measured by the absolute value of the percentage difference between capacity and backlog ((Capacity-Backlog)/Capacity). On average, the gap with the "best" strategy was 80% lower than with the average ballistic strategy.

Second, strengthening the link between delivery delay and price (b3) improves profitability by rationing capacity by price instead of by availability. When there is inadequate capacity, delivery delay increases and market share declines until orders equal

capacity. In this situation, a price increase lowers both market share and delivery delay. Revenue and net income both increase because the higher price more than offsets the smaller market share. Sales are inelastic with respect to price because sales are constrained by capacity and not demand. In addition, additional revenue increases marketing spending and recaptures part of the lost market share.

Third, the lack of outcome feedback in the ballistic rule makes it highly sensitive to errors in anticipating demand. With the ballistic rule, target capacity is based primarily on the initial expectation of market demand. An incorrect a priori forecast of market demand causes a large decrease in performance because outcome feedback is not used to keep capacity and demand in balance.

To illustrate the fragility of the ballistic rule relative to the "best" rule, a simulation experiment was conducted by running two groups of 250 simulations. Fifty simulations were run for each scenario. The first group of simulations was generated by the average parameter values for the ballistic rule. A random component with a zero mean and a standard deviation of .825 was added to the parameter a0, which is the constant in the estimated relationship for target capacity. The standard deviation was computed from the sample of trials that used the ballistic strategy. The second group of simulations used the "best" parameters with the identical random component added to a0.

As expected, randomness has a larger impact on the average ballistic strategy than the "best" strategy. Table nine shows that average performance fell by 38% when randomness is added to the average ballistic parameters and by 18% with the "best" parameters. The theory of feedback systems (Forrester 1961, Richardson 1981) explains that closed feedback loops are essential for correcting deviations between actual and target values of system variables. For example, assume that the initial demand expectation is too high and that too much capacity is added. Target capacity will be pushed closer to actual demand if the size of the order backlog relative to capacity has a

significant weight (a3 is large). This feedback loop negates errors caused by inaccurate demand expectations. In the ballistic strategy, the error-correcting feedback is not present. The lack of feedback causes a significant degradation of performance when there are errors in decision making.

		Table 9	
	No Randomness	Randomness	Ratio
Average Ballistic	$3,550 (millions)	$2,235	62%
Best Parameters	$4,391	$3,604	82%

In addition, the fragility of the ballistic strategy was demonstrated by the subjects' actual performance in the simulation. Because the ballistic strategy is unresponsive to changes in the market, its performance should be much more sensitive to gaps between initial expected demand and actual demand than the adaptive strategy. It was expected that both very high and very low levels of the initial demand expectation (ln(S*De)) would generate poor performance because more extreme values would create large gaps between capacity and actual demand. This non-linear relationship was approximated by a quadratic function form which accounts for the squared term in equation nine. It was hypothesized that the shape of non-linear relationship between the initial demand expectation (ln(S*De)) and performance (POTRAT) could depend on the strategy that was chosen and on the treatment parameters. As discussed above, there was a strong prior belief that performance would be more sensitive to variations in the initial demand expectation when the ballistic strategy was employed.

To test this hypothesis, equation 9 was estimated to account for the impact of the initial demand expectation (ln(S*De)) on the performance of the two basic strategies. The initial demand expectation, ln(S*De), is related to the estimated decision rule

parameters a0 and a1 by $ao/(1-a1)=\ln(S^*De)$. This formula was used to calculate the values of $\ln(S^*De)$ that were used in equation 9.

9. POTRAT=CONSTANT+TRIAL + STRATEGY
+ STRATEGY*LWOM + STRATEGY*LRENT +
$\ln(S^*De) + \ln(S^*De)^2$+
LWOM*$\ln(S^*De)$ + LRENT*$\ln(S^*De)$ +
LWOM*$\ln(S^*De)^2$ +LRENT*$\ln(S^*De)^2$ +
STRATEGY*$\ln(S^*De)$ + STRATEGY*$\ln(S^*De)^2$.

The results of estimating equation 9 are shown in table 10. As expected, the interactions between the choice of strategy (STRATEGY) and the initial demand expectation ($\ln(S^*De)$ and $(\ln(S^*De)^2)$ are highly significant. As discussed below, the interaction effect is consistent with the hypothesis that the ballistic strategy is more sensitive to variation in the initial demand expectation. The interaction between the rate of reentry into the market (LRENT) and both initial expectation terms also are highly significant which is consistent with the fact that the rate of reentry controlled the equilibrium level of demand.

Table 10

	1	2	3	4	5	Significance
Trial	-.176	-.072	.045	.099	.104	.000
Strategy	-9.711	4.933	4.778			.002
LWOM	-.077					.962
LRENT	1.694					.007
Strategy* LWOM	.021	-.132	.11			.014
Strategy* LRENT	-.041	.060	-.019			.391
ln(S*De)	.909					.000
$\ln(S*De)^2$	-.035					.000
Strategy* ln(S*De)	1.466	-.77	-.69			.001
Strategy* $\ln(S*De)^2$	-.054	.029	.025			.001
LWOM* ln(S*De)	.022					.872
LRENT* ln(S*De)	-.326					.002
LWOM* $\ln(S*De)^2$	-.002					.76
LRENT* $\ln(S*De)^2$	.017					.000

The average relationships between the initial demand expectation (ln(S*De)) and performance (POTRAT) for the ballistic and adaptive strategies calculated from the estimated coefficients of equation nine, are graphed in figure 4. The curve for the ballistic strategy has a much sharper peak than the curve for the adaptive strategy which implies that the performance of the ballistic strategy was much more sensitive to errors in

anticipating equilibrium demand. With the ballistic strategy, a 10% deviation from the optimum value of ln(S*De) produces a 28% decrease in performance. With the adaptive strategy, a 10% deviation from the optimum value of ln(S*De) produces a 5% decrease in performance.

The curves shown in figure 4 have different shapes because the adaptive strategy uses outcome feedback. With the adaptive strategy, the expected level of equilibrium demand has little weight in determining target capacity and forecasting errors would be offset by feedback about the actual level of market demand. The ballistic strategy does not use much outcome feedback which makes its performance entirely dependent on the accuracy of the initial demand forecast.

Figure Four

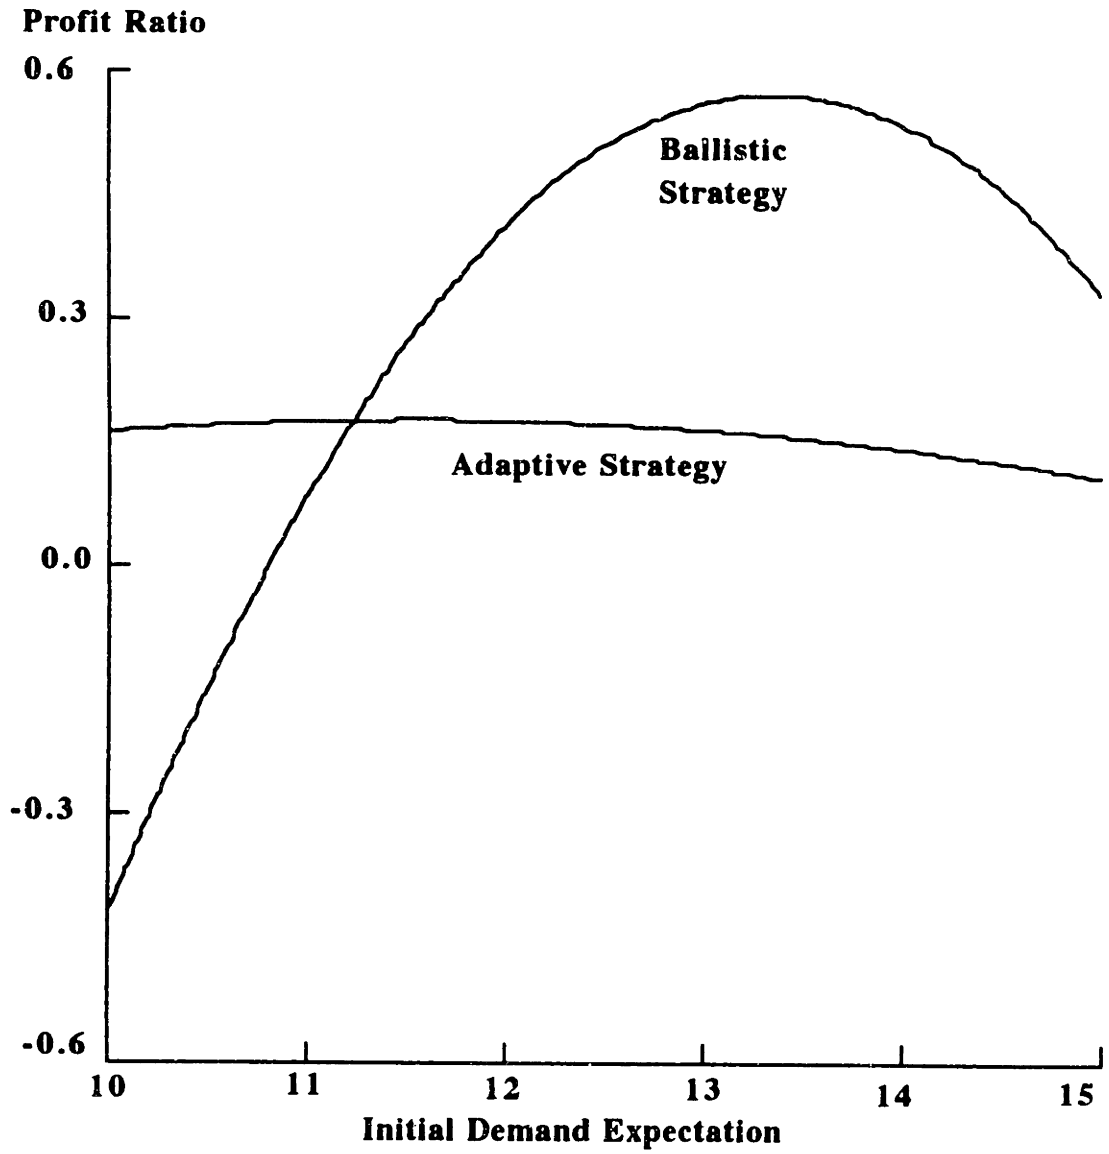
Profit Ratio
0.6
0.3
0.0
-0.3
-0.6
Ballistic
Strategy
Adaptive Strategy
10
11
12
13
14
15
Initial Demand Expectation

For the ballistic strategy, there were significant deviations away from the optimum level of the initial demand expectation. The optimum value of ln(S*De) for the ballistic strategy was calculated for each scenario. The average absolute deviation of ln(S*De) from the optimum for the 228 trials using the ballistic strategy was 1.03 with 47% of the deviations being greater than one.

The ineffectiveness of the ballistic strategy can be seen by analyzing its performance in the individual scenarios. Closer examination revealed that the effect of the ballistic strategy was not consistent across the five scenarios. Table 11 shows the results of estimating a regression for each scenario with relative performance as the dependent variable and trial number and strategy as the independent variables. The regression tested whether or not the strategy had a significant effect on performance for a given scenario. The strategy variable is highly significant for scenarios three and four and insignificant, by a large margin, for the remaining scenarios. The ballistic strategy does not significantly improve average performance in scenarios one, two, and five, which are 60% of the trials.

	Table 11	
	Significance Level	
Dependent Variable:		
Scenario	Trial	Strategy
1	.157	.383
2	.001	.838
3	.00	.000
4	.552	.007
5	.192	.716

Why does the ballistic strategy work for scenarios three and four and not for the others? The answer to this question requires one additional piece of data. Table 12 shows the results of a regression in which ln(S*De) is the dependent variable and trial number and the treatment parameters are the independent variables. The sample includes only those trials (N=228) that used the ballistic strategy. The estimated parameters show that ln(S*De) is related to both LWOM and LRENT. As would be expected, ln(S*De) is positively related to LRENT because the subjects did adjust to the actual equilibrium level of demand.

	Table 12	
Dependent Variable: ln(S*De)		
	Coefficient	**Significance**
Trial	-	.450
WOM	-.236	.012
RENT	.596	.000

The negative relationship between LWOM and ln(S*De) is unexpected and somewhat puzzling. In the simulations, LWOM affects both the growth rate and the equilibrium level of demand. A larger LWOM generates a much larger growth rate and a small increase in equilibrium demand. Apparently, the subjects "learned" that a faster growth rate results in a lower equilibrium level of demand which is exactly the opposite of the relationship actually present in the simulation. One explanation for this result is that, based on experience with previous trials, the subjects perceived that rapid growth signaled "trouble ahead". The expectation of a future collapse in demand caused them to be conservative and resulted in lower values of ln(S*De). Conversely, slower growth may have led the subjects to be more aggressive and produced larger values of ln(S*De).

Many of the subjects' diaries of the experiment discussed the need for a conservative strategy. There was definitely an understanding that word of mouth communication and advertising spending drove growth and that eventually the growth would come to an end. As one subject wrote,

> "A robust and effective strategy would be to aggressively ramp up and down capacity, regarding closely the market growth. There are positive feedback loops which suggest rapid growth , but bringing on eventual collapse. Eventually collapse requires the player to be able to respond quickly to the reduced demand for capacity. A conservative strategy seems to play best."

The inverse relationship between LWOM and ln(S*De) provides an explanation of why the ballistic strategy improved performance in scenarios three and four but not in the others. In scenario two, LWOM and LRENT were both large, equilibrium demand was high, but the inverse relationship between LWOM and ln(S*De) generated an average level of capacity that was only 34% of ***market demand***. The low level of capacity reduced market share and offset the performance gains from the ballistic strategy. In the ballistic strategy, the subjects did not increase price to compensate for capacity shortages so that the loss of market share created a large loss of revenue. In scenario four, LWOM

was large and LRENT was small. The average level of capacity was 53% of market demand which was reasonable and ***did not*** offset the performance gains from the ballistic strategy. In scenario three, the combination of LWOM and LRENT produced average capacity equal to 52% of market demand and, again, did not eliminate the improvement from the ballistic strategy. Scenario five generated capacity equal to 80% of demand which created excess capacity and large losses that could not be offset by the other advantages of the ballistic strategy.

It might be expected that the subjects would have corrected large gaps between initial expected demand and actual demand. This would not have been the case if subjects lowered price to increase orders and utilize excess capacity. Lower price could have masked the excess capacity while causing a direct reduction in revenue and profitability.

If subjects eliminated excess capacity by reducing average price, there should be a negative relationship between the parameter that determines average price, b0, and the target capacity implied by the initial demand expectation, ln(S*De), where ln(S*De) is the natural log of the product of expected market share and the initial expected level of demand. This hypothesis was tested by regressing b0 against ln(S*De), LWOM and LRENT for the 228 trials that used the ballistic strategy. As shown in table 13, the expected negative relationship between ln(S*De) and b0 is highly significant.

	Table 13	
Dependent Variable: b0		
	Coefficient	Significance
WOM	.067	.67
RENT	.310	.065
ln(S*De)	-.347	.002

To summarize, analysis of the individual scenarios reveals a more complex and idiosyncratic relationship between the choice of strategy and performance. The subjects based their expectation of equilibrium demand, in part, on the market growth rate which, in two of the scenarios, caused them to be overly conservative or aggressive. In scenarios three and four, where equilibrium demand matched the predictions of the subjects' mental models, the ballistic strategy performed relatively well. In scenarios two and five, where there was a poor match, there were no gains in relative performance.

For the subjects to make effective use of outcome feedback, they need a reasonable mental model of the system. A good mental model is essential for deciding how to respond to the available cues. For example, in the experiment, to effectively use outcome feedback about the supply/demand balance, the DM must be able to anticipate how a change in target capacity or price will affect sales and profitability.

Especially in the early trials, the subjects found the market dynamics difficult to understand and were unable to create effective mental models of the task. This confusion was reflected in the comments, which were written before each trial, that the subjects wrote for their assignments. Some typical comments follow.

1. Written after the fifth trial,

"If there is a way to succeed, I certainly have not found it. The final cumulative loss was approximately $200 MM. The technique of undercutting the competitor was a failure and my several attempts to regain profitability have not worked."

2. Written after the fourth trial.

> "My scheme has completely failed in this trial. It was to optimistic to think that we can apply the optimal point from a static point of view into a dynamic behavior."

3. Written after the fifth trial,

> "My strategy changed over the course of the game. At first, I just completely missed the boat and had no clue what was going on. The final strategy was to pay attention to nothing other that orders and the competitor's price. Target capacity was decided by looking at orders and ramping up in the first few quarters to stay with it. After about 8 quarters, a target capacity was found that looked as though it would be below the peak and in the ball park where demand ended up. In order to do this, the peak had to be completely ignored and the assumption had to be made that a bust would occur at some point."

III. Discussion

Implications for Dynamic Decision Making Research

The behavioral decision making literature has documented a long list of dysfunctional heuristics that people use in static decision making tasks. Relatively little work has been done on dynamic decision making and there is limited evidence concerning the effectiveness of the heuristics used in dynamic tasks.

Sterman has argued that DMs systematically misperceive feedback relationships in dynamic environments (Sterman 1989a, 1989b). This misperception causes DMs to select decision rules that are inappropriate to the feedback structure of the environment. The present study extended the misperception hypothesis by varying the strength of two important feedback loops as experimental treatments. The results of the present experiment are consistent with the misperception hypothesis. In the early trials, the subjects' performance relative to the baselines consistently varied with the intensity of the two feedback loops. Most of the subjects performed relatively well when the environment was made smooth and simple by weakening the loops. Performance

systematically deteriorated when the feedback loops were strengthened. The subjects' decision rules were well matched to the simple environment and poorly matched to the more complex environment.

The present study also extended the misperception of feedback hypothesis by allowing for the possibility that the subjects would learn improved rules. Performance did improve with experience as the subjects better understood the surface characteristics of the simulation. However, performance remained relatively poor in the strong feedback scenarios, most of the subjects left "money on the table" by failing to integrate the appropriate outcome feedback into their strategies. Indeed, the use of outcome feedback could have significantly improved average performance and made the application of the decision rule much more robust against forecasting errors. In this setting, the most effective strategy was to deal with complex feedback structure in the environment by carefully choosing outcome feedback for the decision rules. In effect, the best plan was "to fight feedback with feedback."

The results of this experiment confirm Sterman's speculation that,

> "It is difficult for experience alone to overcome misperceptions of feedback structure. It is well known that imperfect outcome feedback, noise in the environment, and the uniqueness of many decisions may degrade the learning process. Here [in Sterman's early experiments] outcome feedback is perfect and immediate and subjects make numerous successive choices. One may speculate that it is the subjects' dynamically deficient mental models themselves which hinder learning by focusing attention on inappropriate variables and relationships. Thus, misperceptions of feedback may be self-perpetuating (Sterman 1989b)."

In this experiment, as in Sterman 1989b, feedback was immediate, the subjects made numerous successive choices, and still the misperceptions of feedback did not improve significantly with experience. Even with complete information about the structure of the system, perfect outcome feedback and 200 trials of experience, the subjects continued to use rules that were poorly matched to the environment.

The results of the present study also are consistent with previous work on outcome irrelevant learning structures (OILS) (Einhorn 1980). An OILS occurs when outcome feedback does not produce improved decision making strategies (Einhorn 1980). In the experiment, outcome feedback did not produce robust decision rules, and worse, led the subjects to use of even less within-trial feedback. Because the ballistic strategy worked reasonably well, the subjects never learned that the appropriate use of within-trial feedback could significantly improve performance. In addition, because the ballistic strategy used no outcome feedback, it foreclosed the possibility that experimentation would reveal better decision rules.

Implications for Growth Management: Fighting Feedback with Feedback

The results of the simulation experiments have some implications for the management of growth situations. In the simulations, the subjects have a great deal of difficulty effectively using outcome feedback. Over the course of the five trials, they move toward a ballistic strategy that uses less, rather than more, within-trial outcome feedback. The reduced use of outcome feedback occurs despite the fact that the proper use of such feedback could improve performance significantly. To the extent that the results of the simulation can be generalized to real situations, the implication is that performance in real situations also could be improved by better use of outcome feedback.

Specifically, the subjects make inadequate use of outcome feedback in balancing capacity and demand. In general, the subjects were insufficiently responsive the level and growth of demand and to the supply/demand balance. The subjects failed to adjust price in response to supply/demand imbalances and continued to cut price during the growth phase even though they could not meet existing demand.

The learning that did take place in the game would be much less successful in the real world than in the simulation. By design, the game simplified the environment by holding the potential market constant across scenarios. Steady-state replacement demand ranged from 1.5 to 5 million units which is much less variation than would occur in the real world across successive generations and types of products. The relatively low variation enabled the subjects to make reasonable guesses about the ultimate size of the market. In real situations, previous experience from other markets would be much less useful and no manager could ignore market feedback in making capacity decisions.

The case studies described in chapter two provides some evidence that managers in real situations find it difficult to understand and use outcome feedback. In the Tensor case, Jay Monroe the CEO, remarks that during the growth phase, Tensor could not adequately interpret the sales data (outcome feedback) and that they found it useless in determining production and investment (Salter 1965). This implies that the production and investment decisions must have been based on something other than direct feedback from the market, which is the important characteristic of a ballistic strategy.

In addition, a recent account of Atari's adventure in the video game market suggests that Atari's management ignore numerous cues that the video game market was saturating and that sales would soon decline significantly (Clurman 1992). To an extent, Atari ignored the signals of impending market decline because the CEO of Warner Communication (Atari's parent company), Steven Ross, was unwilling to accept bad news. However, if Atari management had had a more sophisticated shared mental model of the market they might have recognized that declining sales of the existing product generation were not evidence of managerial incompetence but rather the inevitable consequence of the growth and saturation feedback loops. A better mental model might have facilitated communication between managers by producing a more insightful interpretation of the sales data. As described in the case study in chapter two, it is likely

that Nintendo, which seemed to have a better understanding of the underlying market dynamics, would have interpreted declining sales very differently.

To effectively use outcome feedback, the subjects must have a reasonable mental model of the system they are attempting to control. In the simulation, as in real-life, the indiscriminate use of outcome feedback does not automatically produce good outcomes. This is evidenced by the poor performance in early trials that was generated by the tracking strategy which did attempt to use outcome feedback. A good mental model is required for selecting the weights on the relationships that connect the available cues to the decision variables.

In the simulations, the subjects were unable to build good mental models of the situation and resorted to a "ballistic" strategy that mostly ignored the cues generated within a specific trial. On the average, the ballistic strategy performed better than the tracking strategy, but still left a lot of money "on the table." The ballistic strategy was perceived as the best alternative available because the subjects' mental models of the situation remained poor and outcome feedback could not be employed effectively. In addition, the adoption of the ballistic strategy eliminated the possibility that a subject could learn by experimenting with alternative uses of within-trial outcome feedback.

The real-life analogy to the ballistic strategy would be for a DM to base decisions on vivid previous experiences instead of effectively using information from the situation at hand. For example, consider a manager who had been traumatized by a previous market collapse and had not acquired much insight into the dynamic mechanisms that generate market dynamics. In this situation, it is plausible that a DM, who does not have a good mental model of the process, could ignore outcome feedback concerning current market conditions and refuse to increase capacity even if it was clearly warranted. Refusing to change capacity is just one of many idiosyncratic decision rules that could be incorrectly

"learned" from experience and allowed to persist because outcome feedback was no longer utilized.

The concept of "outcome irrelevant learning structures" applies to real growth situations as well as the simulation. Despite the difficulty of building good mental models, some managers do make profits in growth situations. The fact that relatively bad rules can still be profitable is not surprising, even in the simulation, the subjects earn a profit in 83% of the trials. Managers who are reasonably successful may not alter their decision rules until competition and a threat to survival forces them to do so. Since all competitors have difficulty creating mental models of the environment, many rules that are far from optimal can coexist. As Einhorn states,

> "If people learn and continue to use poor rules, doesn't this contradict the evolutionary concept of survival of the fittest? I take this question to mean that those who use bad rules should be less likely to survive than those who use better rules. However, the use of better rules can still be quite removed from the use of optimal rules. The concept of most fit involves a relative ordering whereas optimality implies some absolute level. Therefore, the fact that suboptimal rules are maintained in the face of experience is not contradicted by Darwinian theory. Perhaps the most succinct way of putting this is to quote Erasmus: "In the land of the blind, the one-eyed man is king." (Einhorn 1980)

The above discussion suggests the possibility that a design process that "fights feedback with feedback" could improve performance significantly. The misperception of feedback (MOF) hypothesis holds that interacting feedback loops, delays, and non-linearities make it difficult for DMs to build the mental models which are a prerequisite for effectively using outcome feedback to control dynamic systems. In addition, the experiments discussed above suggest that without effective mental models, DMs may resort to ballistic strategies that are based on experiences generalized from other situations and ignore relevant outcome feedback.

A possible antidote to the complexity introduced by the MOF is the use of simulation and gaming to teach the effective use of outcome feedback. Management games and simulations have long been used to supplement lectures and games. However, despite the fact that many of these simulations are highly complex, students are usually expected to learn from outcome feedback over the course of a single trial. Not surprisingly, evidence on the effectiveness of these simulations is mixed (Graham et al. 1992).

The present study suggests a paradox in the use of business games as a tool to enable learning and improve performance. Feedback processes and time delays are an intrinsic part of all real life situations and they must be included in any meaningful simulation. Yet, the results of the present study provide evidence that the feedback processes and time delays that retard learning in the real system also limit learning in a simulation.

To avoid this learning trap, the objective of the simulation experience should include learning how to use outcome feedback as well as learning from outcome feedback. Initially, the design of such an experience probably must involves active management participation in building a simulation model that represents the essentials of the market situation (Richmond 1990, Sterman 1994). Managerial involvement in the modeling process is essential for getting mental models "out on the table" where they can be carefully scrutinized and debated.

Once created, the simulation model can be helpful in the design of the decision processes that will be used in the real system. The model could be useful in considering questions such as,

1. What information is required to control the system?
2. Do our current information systems provide the required information?
3. What weights should be applied to the various streams of information in making specific decisions?

The concept of putting explicit weights on information cues is foreign to most managers. However, as discussed in the literature review this approach has a long and distinguished history in both the bootstrapping literature and in system dynamics. It is possible to envision an enhanced bootstrapping process in which teams of managers play a simulation game that captures the essence of their own situations. In such a game, it would be possible to determine which teams are most successful in various scenarios and attempt to estimate their decision rules. While it is not clear whether the estimated decision rules could actually be used to make automated choices, the relative magnitude of the weights could provide insight into successful information processing strategies. For example, in the present study we found that the subjects didn't put enough weight on supply/demand imbalances in determining price. This finding could not be used to recommend that the weight on the supply/demand balance should be .15, but such a finding could prompt a useful conversation about the role of the supply/demand balance in setting price. This type of conversation can be viewed as a process for explicitly designing the use of outcome feedback instead of assuming that it automatically will be employed effectively.

There are several serious issues that could limit the effectiveness of this conversation. First, the complexity of real situations makes it difficult to build and analyze an adequate simulation model. Experience in building complex dynamic models has shown that large, detailed models take a long time to build and require extensive validation. Designing robust decision rules requires a deep knowledge of the actual situation as well as an understanding of non-linear feedback systems. For example, designing a real decision rule for capacity investment would require considerable knowledge about the processes and delays involved with capital acquisition as well as an understanding of how these processes can be accurately modeled.

The complexity of model building works against the use of simulation as a tool for designing robust policies. Currently, there are few individuals who combine deep insight into the real situation with good modeling skills and, in practice, outside experts often supply the necessary modeling talent. However, high performing, robust decision rules are often counter-intuitive and are difficult for modelers to explain and for managers to embrace. In such a situation, there is a temptation for decision makers to discredit the model especially if it has been built by outside experts. It is very difficult for effective learning to take place in the face of surprising or counter intuitive findings when managers have only a vague understanding of what is included in a model.

Recently, there have been some promising developments in techniques both for conceptualizing feedback structures and for building simulation models. Techniques for conceptualizing structure include causal loop diagramming (Senge 1990) and hexagon diagrams (Morecroft & Sterman 1994). These techniques enable managers to represent explicitly the structure of their mental models. These techniques are advantageous because they are relatively easy to learn and are accessible to managers who do not have extensive training in model building.

Simulation model building has been made more accessible with the development of graphically based simulation languages (Morecroft & Sterman 1994). In these languages, a model is built by connecting icons on the screen of a personal computer instead of by writing equations. Managers can participate in the model building process by constructing a diagram that can be easily explained and modified. In addition, some of these languages have started to include advanced features for automatically detecting errors in model formulations. Ultimately, these features should make it much easier to establish confidence in the utility of a simulation model.

When dynamic complexity is high, managers must become part of the modeling process. The manager's deep understanding of the real situation should be coupled with

an understanding of how to represent and analyze feedback structure. Recent advances in interactive modeling and the representation of structure have begun to make this possible but a great deal of work remains to be done (Morecroft & Sterman 1994).

References

Angrist, S., "A Fading Rose," *Forbes*, November 7, (1983), 240.

Armstrong, J.S., *Long Range Forecasting*, John Wiley, New York, 1985.

Bakken, B.E., "Learning in Dynamic Simulation Games; Using Performance as a Measure," *Computer-Based Management of Complex Systems*, Proceedings of the 1989 International Conference of the System Dynamics Society, Springer-Verlag, New York, 1989, 309-316.

Bayus, B., "High-Definition Television: Assessing Demand Forecasts for a Next Generation Consumer Durable," *Management Science*, 39, 11 (1993), 1319-1333.

Beihnhocker, E., "Worlds of Wonder (A) and (B)," Case study available from John Sterman, Sloan School of Management, MIT Cambridge MA, 1991

Beinhocker, E., L. Newman and J. Sterman, "A Behavioral Model of Learning Curve Strategy," Working Paper, MIT Sloan School of Management, Cambridge MA, 1993.

Birch D.A., "The Class of 1982," *INC.*, December, (1987), 17-18

Bowman, E.H., "Consistency and Bias in Organizational Decision Making," *Management Science*, 9, 2 (1963), 310-321.

Brandt , R., "Creeping Up on the Mario Bros," *Business Week*, July 30, (1990), 74-76.

Brandt, R., "Sega," *Business Week*, February 21, (1994), 66-74.

Brehmer, B., "Dynamic Decision Making: Human Control of Complex Systems," *Acta Psychologia*, 81 (1992), 211-241.

Brehmer, B., "Strategies in Real Time, Dynamic Decision Making," In R. Hogarth (Ed.), *Insights in Decision Making*, University of Chicago Press, Chicago, 1990, 262-279.

Brehmer, B., "Systems Design and the Psychology of Complex Systems," In J. Rasmussen and P. Zunde (Eds.), *Empirical Foundations of Information and Software Science III*, Plenum Publishing, New York, 1987.

Broadbent, D., and B. Aston, "Human Control of a Simulated Economic System," *Ergonomics*, 21, 12 (1978), 1035-1043.

Business Week, "A Smoother Trail for Snowmobile Makers," December 16, (1976).

Business Week, "Black & Decker Saws Off a Loser," November 8, (1982), 45-46.

Camerer, C., "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence," *American Economic Review*, 77, 5 (1987), 981-987.

Clifford, D, and R. Cavanagh, *The Winning Performance*, Bantam Books, New York, 1985.
Clurman, R., *To the End of Time*, Simon & Schuster, New York, 1992, 114-121.

Cowan, A., "At Last the Game May be Up for Nintendo," *New York Times*, November 25, (1990), 15.

Diehl, E. and J. Sterman, "Effects of Feedback Complexity on Dynamic Decision Making," System Dynamics Group Working Paper D-4401-1, Sloan School of Management, Massachusetts Institute of Technology, March (1994).

Dörner, D., "On the Problems People Have in Dealing with Complexity," *Simulation and Games*, 11 (1980).

Ebert, R.J., "Human Control of a Two-Variable Decision System," *Organizational Behavior and Human Performance*, 7 (1972), 237-264.

EIA Marketing Services, *Electronic Market Data Book*, Electronic Industries Association, 1984

EIA Marketing Services, *Electronic Market Data Book*, Electronic Industries Association, 1989.

EIA Marketing Services, *Electronic Market Data Book*, Electronic Industries Association, 1993.

Einhorn, H.J., and R. Hogarth, "Behavioral Decision Theory: Process of Judgement and Choice," *Annual Review of Psychology*, 32 (1981), 53-88.

Einhorn, H.J., "Learning from Experience and Suboptimal Rules in Decision Making," In T. Wallsten (Ed.) , *Cognitive Processes in Choice and Decision Behavior*, Erlbaum, Hillsdale, N.J., 1980.

Forbes, "Emerging Growth Companies: Why So Many Peak So Early, January 28, (1985), 69-75.

Forrester, J., *Industrial Dynamics*, MIT Press, Cambridge MA, 1961.

Forrester, J.W., "Counterintuitive Behavior of Social Systems," *Technology Review*, 73 (1971), 52-68.

Forrester, J.W., "Market Growth as Influenced by Capital Investment," *Industrial Management Review*, 9, 2 (1968), 83-105.

Funke, J., "Solving Complex Problems: Exploration and Control of Complex Systems," In R Sternberg and P. Frensch (Eds), *Complex Problem Solving: Principles and Mechanisms*, Lawerence Erlbaum Associates, Hillsdale, NJ, 1991.

Funke, J., "Using Simulation to Study Complex Problem Solving," *Simulation and Games*, 19, 3 (1988).

Graham, A.K., J.D. Morecroft, P.M. Senge, and J. Sterman, "Model Supported Case Studies for Management Education," *European Journal of Operational Research*, 59, 1 (1992), 151-166
Harrigan, K., "An Application of Clustering for Strategic Group Analysis," *Strategic Management Journal*, 6, 1 (1985), 55-74.

Hartigan, J., *Clustering Algorithms*, John Wiley, New York, 1975.

Hill, G., "Nintendo-Atari Zapping Contest Goes to Washington," *Wall Street Journal*, December 8, (1989), B1-B6.

Hogarth, R., and M.W. Reder, *Rational Choice: The Contrast between Economics and Psychology*, University of Chicago Press, Chicago, 1987

Hogarth, R., and Makridakis, S., "Forecasting and Planning: An Evaluation," *Management Science*, 27 (1981), 115-138.

Hogarth, R., and Makridakis, S., "The Value of Decision Making in a Complex Environment: An Experimental Approach," *Management Science* , 27, 1 (1981), 93-107.

Hogarth, R., "Beyond Discrete Biases: Functional and Dysfunctional Aspects of Judgemental Heuristics," *Psychological Bulletin*, 90 (1981), 197-217.

ISIA, "U.S. Snowmobile Sales," International Snowmobile Industry Association, Annandale Virginia, 1985.

Kampmann, C. and J. Sterman, "Do Markets Mitigate Misperceptions of Feedback in Dynamic Tasks?" *Proceedings of the 1992 International System Dynamics Conference*, Utrecht, The Netherlands, 1992.
Kleinmuntz, D. and J. Thomas, "The Value of Action and Inference in Dynamic Decision Making," *Organizational Behavior and Human Decision Processes*, 39, 3 (1987), 341-364.

Klepper, S., and E. Graddy, "The Evolution of New Industries and the Determinants of Market Structure," *Rand Journal of Economics*, 21, 1 (1990), 27-44.

Kluwe, R.H., C. Misiak, and O. Ringelband, "Learning to Control a Complex System: The Effects of System Characteristics on System Control," *Research Report No. 12*, Institut fur Kognitionsforschung, Universitat der Bundeswehr Hamburg, (1986).

Leamer, E. E., "Let's Take the Con Out of Econometrics," *American Economic Review*, 73, 1 (1983), 31-43.

Lilien G., and P. Kotler, *Marketing Models*," Prentice Hall, Englewood Cliffs, N.J., 1992.

Lilien G., and S. Kalish, "Applications of Innovation Diffusion Models in Marketing," In Vijay Mahajan and Yoram Wind (Eds), *Innovation Diffusion Models of New Product Acceptance*, Ballinger, Cambridge MA, 1986.

Mackinnon A.J., and A.J. Wearing, "Complexity and Decision Making," *Behavioral Science*, 25, (1980), 285-298.

MacKinnon, A.J., and A.J. Wearing, "Systems Analysis and Dynamic Decision Making," *Acta Psychologica*, 58, (1985), 159-172.
Mahajan, M , E. Muller, and F. Bass, "New Product Diffusion Models in Marketing: A Review and Directions for Research," *Journal of Marketing*, 54, (1990), 1-26.

Main, J., "A Chipmaker Who Beats the Business Cycle," *Fortune*, December 23, (1985), 114-120.

McCann, J., J. Miller and H. Moskowitz, "Modeling and Testing Dynamic Multivariate Decision Processes," *Organizational Performance and Human Performance*, 14 (1975), 281-303.

McClellan, S., *The Coming Computer Industry Shakeout*, John Wiley, New York, 1984.

Morecroft, J.W., and J. Sterman, *Modeling for Learning*, Productivity Press, Portland OR, 1994.

Moskowitz, H., "Regression Models of Behavior for Managerial Decision Making," *OMEGA*, 2,5 (1974), 677-690.

Neff, R., "The Newest Nintendo Will Take a Slow Boat to America," *Business Week*, July 2, (1990), 46.

Neter, J., W. Wasserman and M. Kutner, *Applied Linear Statistical Models*, Irwin, Homewood Ill., 1985.

Paich, M., "Generic Structures," *System Dynamics Review*, 1 (1985), 126-132.

Pereira, J., "Nintendo is Counting on New Super Game to Rescue U.S. Sales," *Wall Street Journal*, May 10, (1991), 1.

Petre, P., "Jack Tramiel is Back on the Warpath," *Fortune*, March 4, (1985), 46-50.

Pindyck, R., and D. Rubinfeld, *Econometric Models & Economic Forecasts*, McGraw-Hill, New York, 1991.

Plott, C.R., "Rational Choice in Experimental Markets," *The Journal of Business*, 59 (1986), S301-S327

Popper, A., "Worlds of Wonder: From Wall Street Charmer to Chapter 11," *Business Week*, March 21, (1988), 74-75.

Porter, M., *Competitive Strategy*, Free Press, New York, 1980.

Porter, M., *The Chain Saw Industry in 1978*, Harvard Business School Publishing Division, Boston MA., 1981.

Rappaport, A., "A Research Paradigm for Studying Dynamic Decision Behavior," In D. Wendt and C. Vlek (Eds), *Utility, Probability, and Human Decision Making*, Dordrecht, Holland: Reidel, (1975) 349-374.

Remus, W.E., "Testing Bowman's Managerial Coefficient Theory Using a Competitive Gaming Environment," *Management Science*, 24, 8 (1978) 827-835.
Richardson, G., and A. Pugh, *Introduction to System Dynamics Modeling with Dynamo*, Productivity Press, Cambridge MA, 1981.

Richmond, B., *The Strategic Forum: From Vision to Operating Policies and Back Again*, High Performance Systems, 13 Dartmouth College Highway, Lyme NH.

Roman, M., "Oopsie-Daisy, Guess Who's Tops in Toyland Now," *Business Week*, December 18, (1989), 102.

Salter, M., *Fuqua Industries*, Harvard Business School Publishing Division, Boston MA., 1974.

Salter, M., *Tensor Corporation Case*, Harvard Business School Publishing Division, Boston MA., 1969.

Sargent, T.J., *Macroeconomic Theory*, Academic Press, New York, 1979.

Schlesinger, J., "Electronics Industry in Japan Hits Limits After Spectacular Rise," *Wall Street Journal*, April 28, (1992), 1.

Senge, P., "A System Dynamics Approach to Investment Function Formulation and Testing," *Socioeconomic Planning Sciences*, 14 (1980), 269-280.

Senge, P., *The Fifth Discipline: The Art and Practice of the Learning Organization*, Doubleday, New York, 1990.

Simon, H.A., *Models of Bounded Rationality*, The MIT Press, Cambridge, MA, 1982.

Simon, H.A., "Rationality in Psychology and Economics," *The Journal of Business*, 59 (1986), S209-S224.

Slovic, P., B. Fischhoff and S. Lichtenstein, "Behavioral Decision Theory," *Annual Review of Psychology*, 28 (1977), 1-39

Smith, G., "Emerging Growth Stocks-Why So Many Peak So Early,"

Smith, V., "Experimental Methods in the Political Economy of Exchange," *Science*, 234, (1986), 167-173.

Smith, V., "Experimental Microeconomic Systems," *American Economic Review*, 72, 5 (1982), 924-952.

Smith, V., G. Suchanek and A. Williams, "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 5 (1988), 1119-1152..

Stalk, G., and T. Hout, *Competing Against Time*, Free Press, New York, 1990.

Sterman, J., "Expectation Formation in Behavioral Simulation Models," *Behavioral Science*, 32 (1987), 190-211.

Sterman, J., "Learning In and About Complex Systems," *System Dynamics Review*, February (1994).

Sterman, J., "Misperceptions of Feedback in Dynamic Decision Making," *Organizational Behavior and Human Decision Processes*, 43, 3 (1989a), 301-335.

Sterman, J., "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment, " *Management Science*, 35, 3 (1989b), 321-339.

Sterman, J.D., "People Express Management Flight Simulator," Simulation Game Briefing Book, and Simulator Guide, Available from Author, MIT Sloan School of Management, Cambridge, MA, 1988.

Tellis, G., "The Price Elasticity of Selective Demand: A Meta-analysis of Econometric Models of Sales," *Journal of Marketing Research*, 25 1988, 331-341.

Urban, G., and J. Hauser, *Design and Marketing of New Products*, Prentice Hall, Englewood Cliffs N.J., 1980, 408.

Wagenaar, W., and H. Timmers, "The Pond-and-Duckweed Problem: Three Experiments in the Misperception of Exponential Growth," *Acta Psychologica*, 43 (1979) 239-251.

Wasson, C.R., *Dynamic Competitive Strategy and Product Life Cycles*, Challenge Books, St. Charles Ill., 1974.

Willard, G.E., and A.C. Cooper, "Survivors of Industry Shake-outs," *Strategic Management Journal*, 6, 4 (1985), 299-318.