EFFECTS OF FEEDBACK STRUCTURE ON DYNAMIC DECISION MAKING

by

Ernst-Walter Diehl

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

June 1992

Signature of Author..
Sloan School of Management

Certified by..
John D. Sterman
Assoc. Professor of Management Science
Thesis Supervisor

Accepted by..
James B. Orlin, Chairman
Ph.D. Committee

EFFECTS OF FEEDBACK STRUCTURE ON DYNAMIC DECISION MAKING

by

ERNST-WALTER DIEHL

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

ABSTRACT

Prior research shows that decision making effectiveness varies widely from situtation to situation. In dynamic decision making research, in particular, some experiments suggest human performance is close to optimal, while others show dramatic, persistent dysfunction. There has been no theoretical framework to explain differences in dynamic decision making performance as a function of the task environment.

I propose that feedback theory can provide a taxonomy that is able to relate complexity of decision situations to the likelihood of dysfunctionality of human decision strategies. The research challenge is to determine what particular characteristics in feedback structures lead people to perform poorly or well.

An experimental study was carried out to determine subjects' performance in a stock-adjustment task. Subjects were asked to be production managers and to control a stock of inventory in the face of varying sales. A Latin Square design was used with two treatment conditions: varying delay in production and varying side effect of production on sales. Simulations of two benchmark decision rules were performed to compare subjects' performance against: no-control case and optimum-control case.

Subjects performed worse than the optimum-control case across conditions. Subjects outperformed the no-control case in the low delay, low side effect condition. With increasing delay and side effect, the performance difference decreased. In the high delay, high side effect condition, the majority of subjects was outperformed by the naive no-control rule.

Subjects persistently undercontrolled the system and did not exert as much control as suggested by the optimum-control rule. Undercontrol increased with increasing delay and increasing side effects; showing that subjects' understanding of complex feedback settings declines as delays between cause and effect increase, and as actions have stronger side effects. Performance did not improve after the third trial. Little indications were found of active experimentation: the need to control seemed to override the ability to learn.

Thesis Committee:

John Sterman (Assoc. Professor of Management Science, M.I.T. , chair)
John Carroll (Professor of Behavorial and Policy Sciences, M.I.T.)
Don Kleinmuntz (Assoc. Professor of Accountancy and Decision Research,
University of Illinois at Urbana-Champaign)

Acknowledgments

While working on this dissertation I had the good fortune to meet many talented and insightful mentors, colleagues and friends. This thesis would not have been possible without them and I want to acknowledge and thank them for their help.

Foremost I would like to thank John Sterman, who encouraged me to come to MIT and who guided me throughout the entire thesis process. My work is very much influenced by John's ideas and his pioneering work in bringing system dynamics and behavioral decision theory together. John Carroll and Don Kleinmuntz, the two other members of my thesis committee, also deserve much thanks. My thesis profited significantly from their detailed comments and advice. I owe a deep sense of gratitude to Jay Forrester whose writings introduced me to the field of system dynamics and shaped my career decisions.

I profited from many discussions at conferences and seminars, too numerous to mention. I would like to acknowledge Berndt Brehmer and Dietrich Dörner in particular, for their valuable input at decisive moments in my research and for sharing their working papers with me.

I also would like to thank all my current and former colleagues for the supportive atmosphere and academic environment they have created within the System Dynamics Group at MIT: Bent Bakken, Janet Gould-Kreutzer, Alan Graham, Jim Hines, Bill Isaacs, Christian Kampmann, Jenny Kemeny, Daniel Kim, David Kreutzer, John Morecroft, Mark Paich, Peter Senge, and Becky Waring. I got to know all of them not only as scholars but also as friends.

My special thanks go to Peter Senge and Diane Leonard-Senge for the spirit of community they provided; the annual ski weekend is just one of the events that would not have happened without them. Much appreciation goes to Angela Lipinski, Nan Lux, and Eileen Mullen for their reliable support and assistance they always provided.

Finally a big thanks for Michael Vandall for the outstanding support he provided me with. His diligence was outstanding and he never complained when I came up with yet another figure to include into the thesis. Thank you, Michael!

I gratefully acknowledge the financial support of the Studienstiftung des Deutschen Volkes.

To Christine.
Far away, yet always near.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF FIGURES-CONTINUED

LIST OF FIGURES-CONTINUED

LIST OF FIGURES-CONTINUED

LIST OF FIGURES-CONTINUED

LIST OF FIGURES-CONTINUED

1. INTRODUCTION

1.1. Motivation

Relating the complexity of decision situations to the likelihood of dysfunction in human decision strategies is a major research challenge in behavioral decision theory (Hogarth, 1981; Kleinmuntz, 1985). The potential rewards for meeting the challenge are large. The media remind us on a daily basis about the severe consequences of human malfunction under complex decision situations. Despite our best intentions, inner cities decline, national economies experience severe cycles of recession, and nuclear reactors come close to melt-down.

Most of the time, flaws in our decision making play out on an individual level and do not make national headlines, yet their persistence and severity is no less troubling. Behavioral decision theorists have documented numerous, persistent shortcomings in the way we make decisions. The list of recognized decision biases is long and still growing; some common traps are that we tend to make predictive judgments based on variables that we believe to co-vary when in fact they do not (Chapman, 1967), we do not appreciate the influence of chance factors and misinterpret events in favor of our particular view (Brehmer, 1980), planning or forecasting activities give us an illusion of control over future events (Langer, 1975), and we put undue emphasis on memorable events and tend to ignore statistical facts (Bar-Hillel, 1980; Tversky, and Kahneman, 1973).

Visitors from outer space, with behavioral decision literature as their only travel guide, might be inclined to ignore humans and look for a second race on Earth. "Where is the race," they might ponder, "that built the pyramids, made important medical discoveries, designed elaborate social systems, and put a man on the moon?" The apparent discrepancy between positive and negative achievements has puzzled some natives of earth as well (Toda, 1962). The apparent dichotomy between adequate and inadequate decision abilities must be understood better to prevent the undesirable consequences of our decision biases. While we know many potential biases that can lead to bad decisions, we do not know how to recognize a potentially troublesome decision situation when we encounter one.

Although scientists have studied human decision making in such different application domains as fire fighting (Brehmer, 1989), medical diagnostics (Kleinmuntz, and Thomas, 1987), management of welfare agencies (Mackinnon, and Wearing, 1980), and inventory management (Sterman, 1989a), comparisons across domains are difficult and extrapolations to new situations are questionable as long as we have not developed a framework that will allow for comparisons and theory over the entire range of human decision activities.

1.2. A taxonomy for unifying dynamic decision tasks

Enumeration and description of phenomena without a theory that incorporates how task characteristics affect information processing may merely lead to stock piles of data

(Fleishman, Quaintance, and Broedling, 1984). Or as George Miller once described scientific journals, "...catalogs of spare parts for a machine they never build" (Miller, 1956, p. 252). It becomes increasingly more important as more research is generated to connect the various facts through theory. However, a suitable theory can only come about as better descriptions of the task environment evolve. Additionally, there is a need to characterize tasks such that researchers and technologists may compare, communicate, and apply the research findings of the various disciplines (Fleishman et al., 1984). Fleishman (1984, pp. 1-3), in his book on taxonomic development for behavioral scientists, states,

> "There is a need to conceptualize tasks and their characteristics to resolve central problems in the study of human behavior. If we are going to generalize about conditions affecting human performance, it is necessary to consider the properties of tasks as important constructs in psychological research and theory as well as in our conceptions of human work and achievement. Such constructs will help to address many common concerns in basic and applied psychology and to integrate concepts and research in a number of seemingly diverse fields....Lacking an organizing framework, experimentalist and technologist alike find generalization, communication, and application of research findings to be difficult. Both behavioral technologist and scientist struggle to relate their results to those from previous studies, similar situations with which she or he has yet to deal, or to the findings of researcher and technologists working on allied problems."

He further relates that some benefits for the taxonomic development of tasks for scientists include:

> 1) conducting literature reviews
>
> 2) establishing better bases for conducting and reporting research studies to facilitate their comparison
>
> 3) standardizing laboratory methods for studying human performance
>
> 4) generalizing research to new tasks
>
> 5) exposing gaps in knowledge
>
> 6) assisting in theory development. (Fleishman et al., 1984, pp.5-6)

Hence, the development of a taxonomy of tasks, including dynamic decision tasks, would provide the force for greater scientific development and increase our understanding of sequential decision making.

I believe that feedback theory can provide such a taxonomy that is able to relate complexity of decision situations to the likelihood of dysfunctionality of human decision strategies (Sterman, 1989a). The research challenge as stated at the beginning of this thesis can then be reformulated: to determine what particular characteristics in feedback structures lead people to perform poorly or well.

A closed-loop view of the world lies at the heart of feedback theory. An open-loop view assumes a linear relationship between cause and effect, in which behavioral effects depend on causes but not vice versa. Feedback theory closes the loop (Powers, 1973). It argues that whenever effects persist in time (and it is hard to think of any effect that does not meet the persistency criteria), the distinction between effect and cause becomes meaningless: time-persistent effects become the "cause" for further behavioral effects.

Feedback theory views the decision maker as a controller integral to the loop. Decision-making is viewed as goal-directed behavior where the decision maker reacts to disturbances by choosing appropriate actions to bring the system back to a desired state. Figure 1.1. illustrates the concept. The figure reveals the essential control loop which contains five main components. Beginning with the **goal**, the task environment produces some gap between an a priori desired behavior and actual system behavior. The **decision maker** must take actions to reduce the discrepancy. Between the action and the resulting outcomes, the **task environment** intervenes with a varying amount of complexity that must be captured in a framework that will facilitate a greater understanding of the interaction of both the decision maker and the environment. **Action feed-out** and **outcome feed-in** link the decision maker and the task environment.

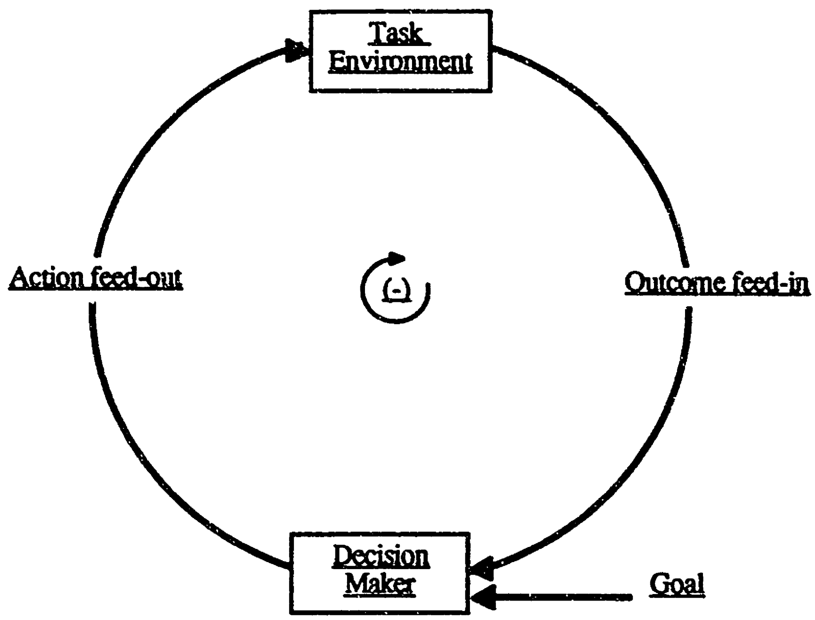

Figure 1.1 The essential control loop

The essence of everyday life and certainly of dynamic decision tasks manifests itself within this framework of closed-loop control. As Powers (1973, p.351) puts it:

> "Feedback is such an all-pervasive and fundamental aspect of behavior that it is as invisible as the air that we breathe. Quite literally it is behavior - we know nothing of our own behavior but the feedback effects of our own outputs. To behave is to control perception."

Given the breadth and importance of the research challenge at hand it comes as no surprise that practitioners in different fields have investigated function and dysfunction of human decision strategies. I will concentrate on four fields from which my research draws: 1) Supervisory control, an application area that is aimed at understanding how humans function as part of human-machine systems and how complex machinery such as nuclear reactors should be constructed to suit human controlling capabilities, 2) Experimental economics, an area that has developed a well-defined experimental protocol and has succeeded in mapping out function and dysfunction of humans in the narrow but important domain of humans as traders in economic markets, 3) Behavioral decision theory, a field concerned with humans as problem solvers and decision makers, and finally 4) Dynamic decision theory, an area concerned with sequential decision making in complex environments.

Where possible, I will link the literature to the essential control-loop framework proposed in this thesis. In particular, I will discuss to what extent the researchers cited are cognizant of the underlying control loop and/or how their research concentrates on and illuminates single elements within the control loop such as the outcome feed-in link or learning processes within the decision maker.

We will see that the need for a taxonomy of tasks is recognized in almost all the research cited; very often the struggle for a taxonomy is phrased as the need to characterize the "complexity" of the task environment. We will conclude the literature review with a summary that lists the dimensions of complexity which have been studied in the different fields and relates them to the feedback framework proposed in this thesis.

1.3. Literature Review

1.3.1 Supervisory control

Understanding and managing the behavior of dynamic systems is the classic domain of **control theory**. The objective of the theory in general terms is to find a control rule so that the actual behavior of the task to be controlled is close to a desired behavior. Once determined, the control rule is usually implemented in physical form (as an automatic controller). As a result, control theory is strictly normative. Whenever possible, it determines what is the optimal control rule.

While the main focus of control theory is automatic control, there exists a small but important literature that focuses on human control. Early work was motivated by the need to understand the human operator's characteristics in order to integrate him effectively into a man-machine system (Crossman, and Cooke, 1974; Sheridan, and Ferrell, 1974). More recent research reflects the fact that humans have increasingly become supervisors of complex automated machinery instead of being operators (Rasmussen, and Rouse, 1981). Automatic control rules typically work only in a range of normal conditions. Human supervisors are still needed to regulate unusual cases and unexpected failures. To understand how humans perform this control is understandably of great interest in such applications as the supervision of nuclear power plants (Lees, 1974; McLeod, 1976).

In designing an automatic control the only bounds considered are whether it is physically feasible to construct a controller with the required specifications. For ensuring optimal human process control and supervisory control, bounds on human decision making have to be considered in addition. Coming from an engineering background the authors in the field address the bounds on human decision making typically in a solution-oriented manner.

Topics that have received particular attention include the effects of information display on performance (Attwood, 1974; Sheridan, and Johannsen, 1976; West, and Clark, 1974), the effects of training on performance (Kragt, and Landeweerd, 1974; Mann, and Hammer, 1986; Moray, and Pam Pajak, 1986; Morris, and Rouse, 1985; Sheperd, Marshall, Turner, and Duncan, 1977) and the performance difference of novices and experts (Bainbridge, Beishon, Hemming, and Splaine, 1974; Mann et al., 1986; McLeod, 1976). Different types of information displays influence the quality of the outcome feed-in link of the feedback control loop, while differences in training and skill levels affect the extent to which goals are learned and information is transformed into action that results in outcomes close to desired goals.

Results of experiments in process control and supervisory control are published in _Ergonomics_, _Human Factors_ and _IEEE Transactions on Systems, Man and Cybernetics_. Many studies in the field are difficult to access or remain unpublished, however. See Edwards and Lees (1972) and Lees (1974) for an overview and Edwards and Lees (1974) for a collection of some better known studies.

1.3.2 Experimental Economics

While studies in supervisory control investigate how subjects attempt to control machinery, experimental economics is interested in learning more about humans as traders of goods and services. In this domain, experimental economics has been able to shed light on the role that the task environment (namely different trading institutions) has on the function or dysfunction of decision makers (namely the speed of markets' convergence toward equilibrium and the nature of the equilibrium reached).

Chamberlin (1948) and Smith (1962, 1965, and 1967) pioneered the field of experimental economics. Smith's point of departure is the idea put forward by Walras (1954) that for equilibrium theory to be a useful concept, one needs to be able to imagine at least one reasonable mechanism that would enable ordinary economic players to achieve equilibrium in actual economic settings. Walras put forward the notion of an auctioneer who calls out prices in an attempt to clear the market. The auctioneer collects all bids and offers for a posted price. No trade takes place until the auctioneer has hit on the equilibrium price where the number of bids equals the number of offers. Instead of Walras' auctioneer mechanism, where trades take place only at the equilibrium price, Smith used experimental settings where individual market participants are allowed to trade at any price they agree on. Smith thus replaced Walras' impersonal market mechanism, where an auctioneer steers the market to the "right" point, with a mechanism where adaptation is done exclusively by market participants making decisions to influence resulting outcomes to their desired goals. In doing so, Smith has expanded the scope of traditional economic equilibrium theory which is mute on the issue of how equilibria might actually be achieved.

Of particular interest for the current study are Smith's more recent papers (Plott, and Smith, 1978; 1976, and 1982), where the authors explain the experimental method and its justification in great detail. The bulk of experimental economics is aimed at testing and exploring the domain of economic theories. Experimenters try to insure external validity or "parallelism" as Smith (1982) calls it by giving particular care to (a) incorporating the essential features of actual economic institutions into the laboratory experiment and (b) inducing predictable preferences by rewarding subjects according to a payoff schedule (Smith, 1976, p. 275).

Investigating the effects of different trading rules, or trading institutions, is a continuing motivation for much of the experimental work. Having early on demonstrated that real markets exist that lead to prices and quantities exchanged as predicted by economic theory (Smith, 1962, and 1965), researchers subsequently started to focus on the question: "How efficient are different trading mechanisms in leading the market to an equilibrium?" (Holt, 1980; Plott et al., 1978; Smith, Williams, Bratton, and Vannoni, 1982; Williams, 1973).

The results indicate that the convergence toward an equilibrium is sensitive to the trading mechanism assumed. The double oral auction institution provides for the most rapid path toward convergence. In a double oral auction market, both buyers and sellers are allowed to cry out bids and offers. By contrast if the rules are changed to a system of posted price, where only buyers or sellers can post prices and prices stay in effect for the entire trading period, convergence to an equilibrium is markedly delayed. If sellers can post "take it or leave it" bids, prices tend to stay above equilibrium for long times; if buyers are the only ones who can post bids, prices stay below equilibrium for many periods (Plott et al., 1978,pp. 146-47; Williams, 1973,p. 110). Constraints on the action feed-out link, such as who can post prices and how frequently prices can be posted, influence decision makers to misinterpret their task environment, resulting in slower movement to equilibrium.

While much of the early work was primarily interested in static efficiency and convergence properties, researchers have only recently begun to focus specifically on the nature of dynamic behavior exhibited in the experiments. (Garner, 1982; Plott, 1982). Or, as Shubik (1979, p. 354) notes:

> "Many different institutions may have the same static efficiency properties, but it is possible that they manifest considerably different dynamic properties. The questions concerning the selection of optimal ... institutions in a fully dynamic context have hardly been asked in a precise form, let alone answered."

We have reviewed two fields which were very concerned with research and conclusions for their particular application domain, namely humans as controllers of machinery and humans as traders of goods and services. In contrast, behavioral decision theorists are interested in how humans perform consistently across applications.

1.3.3 Behavioral decision theory

The goal of behavioral decision theory is to describe the processes and the limitations and strengths of human decision making. People make decisions based on their individual understanding of the world. Cognitive psychologists (Best, 1986) and behavioral decision theorists (Hogarth, 1987) characterize people's thinking processes in terms of their ability to process information. This research paradigm has revealed that decision makers are severely limited information processors with respect to the amount of information available (i.e., outcome feed-in) from their task environment. In order to deal with this amount of information, decision makers adapt by selecting and simplifying whenever possible. This selection and simplification process takes place in four areas of information processing: acquisition, virtually sequential processing, processing capacity, and memory.

The effects of limited information processing have been described by Hogarth (1987,pp. 4-7):

> 1) Perception of information is not comprehensive but selective. We utilize scripts with which to encode or acquire information. We perceive as we believe.

2) Information is processed almost exclusively in a sequential manner. We are limited in the amount of information that we can integrate at one time, but we can integrate information, and usually do, over time. This affords us the opportunity to learn causal relations that are stable and constant and relatively close together in time and space, but in an unstable setting, this type of processing proves faulty, especially with delays in feedback of actions taken.

3) People have limited processing capacity. Instead, they use simple heuristics or rules of thumb to reduce mental effort.

4) We have limited memory capacity. Our memory reconstructs information instead of retaining exact details.

Hence, our ability to cope with a world that is abundant in information is limited, yet we have evolved in such a way that has proved quite efficient. Modern task environments demand that people make decisions in which their limitations prove to be inefficient or costly. For example, a nuclear power plant operator's actions are subject to extreme delays that make controlling the power plant very difficult.

The information processing stages give rise to the many dysfunctions we see in decision making. Many of the important biases were mentioned in the introduction, but they will be repeated here along with the information processing stage from which they emanate and the component of the feedback loop that they affect. Hogarth (1987, pp. 209-215) describes:

a) Acquisition [Outcome feed-in] - These biases involve if and in what way information from memory or the environment becomes salient. We selectively perceive what we are expecting to see even though objective reality is different. The ease of recall of information from memory affects determinations of frequency which is called the availability bias. Availability of cues in environment guides decision making, i.e. the way a problem is set-up can steer the availability and representativeness of information.

b) Processing [Decision maker] - Our ability to process large quantities of information becomes simplified by using an inappropriate strategy, and we are inconsistent appliers of that decision strategy even if situations are identical. For instance, in the face of many pieces of data with which to make a predictive choice, we anchor on a particularly salient one and adjust our decisions from that anchor. Likewise, we ignore statistical facts and base our predictive classifications on qualities of how representative an instance is of a group.

c) Output [Action feed-out] - The manner in which responses are asked to be made influences judgements. We estimate probabilities of occurrence of an event differently under varying experimental and everyday conditions. Another example would be that we have an illusion of control over future events if we merely engage in activities such as planning or forecasting.

d) Feedback- Feedback influences the extent to which we learn the success of our judgements and choices. As feedback is distorted or delayed, learning becomes more difficult. For example, we seek to confirm our hypotheses instead of seeking information that would disconfirm them thereby reenforcing erroneous beliefs. Or we have unwarranted overconfidence in our decisions even if we have only sketchy information. Furthermore, we do not appreciate the influence of chance factors and misinterpret events in favor of our particular view.

e) Interactions of the above biases - All of the above information processing stages may interact and influence one another. For instance, the requirement to make a particular response can direct attention to select information from memory or the environment, perhaps causing the wrong information to be selected or attended to.

The biases of the various information processing stages were mostly discovered in static decision making research. In the early 60's Edwards (1962) and Toda (1962), and recently Hogarth (1981) have all advocated the study of decision making within a more dynamic, but still controlled, setting. These researchers point out that decision making normally takes place in a dynamic context as opposed to a static one. It is, therefore, important to understand to what extent the previously discovered biases need to be modified with respect to dynamic settings.

Edwards (1962) relates that early work in decision making focused on static decision tasks in which a person is confronted by an environment he is supposed to evaluate and come to a possible course of action with the goal of maximizing a payoff. He does this by deciding what payoffs will be associated with various courses of action. However, upon taking a course of action, the person receives his payoff and the world ends. The decision maker is never aware of the consequences of his actions upon the environment (i.e., action feed-out) and, therefore, he neither is required nor needs to learn from the consequences of his actions (i.e., outcome feed-in). This represents a rare decision situation. Many decision situations which confront people in life do provide feedback about consequences, afford learning, and even require learning. While behavioral decision theory provided initial explanations of the basic processes of decision making, dynamic decision theory extended those results to dynamic contexts.

1.3.4 Dynamic decision theory

Maintaining performance in an uncertain and ever changing environment is a task that confronts individuals as well as organizations. Decision makers constantly adjust their actions to cope with changes in the state of a system. The new state of the system influences their next decision.

Hopefully, decision making goes on continually refining itself. Sequential decision making provides the opportunity for the decision maker to utilize outcome feed-in making his or her next decision. Nonetheless, sequential decision making may be subject to the same biases as static decision making. Therefore, one goal of dynamic decision theory is to determine how findings about information processing biases need to be modified in sequential decision making tasks.

Hogarth (1981) makes a convincing argument that the heuristics operating in a static decision task, which lead to errors and biases, may be quite functional in a dynamic situation if the decision maker is considered to be a learning organism whose goal is to learn about the environment and apply creativity in order to induce control over the system. Thus, what we normally see as a deficiency in discrete tasks might be relatively functional in a dynamic setting. Hogarth describes this problem quite well by pointing out the fact that a person's commitment to any one choice can be reduced if the environment allows for correction of errors or learning along the way.

It is important to note that although people may seem more functional in dynamic settings versus discrete ones, they remain limited information processors with fallible mental representations, and they make errors. Often, people are called upon to make decisions in situations more complex than their models afford without much opportunity for learning or exploration. Hogarth (1987, p. 226) relates,

> "In short, the efficacy of feedback depends on the ability to interpret it, and-unless one

has a well developed causal schema for a given task- use of cues-to-causality... such as co-variation and contiguity in time and space may be quite misleading in complex situations."

The point is that decision making over time may help by allowing learning, but it may hurt because of additional complexity to be managed. Tversky (1986, pp. 274-5) adds:

"Effective learning takes place only under certain conditions: it requires accurate and immediate feedback about the relation between the situational conditions and the appropriate response. The necessary feedback is often lacking for the decisions made by managers, entrepreneurs, and politicians because (i) outcomes are commonly delayed and not easily attributable to a particular action; (ii) variability in the environment degrades the reliability of the feedback, especially where outcomes of low probability are involved; (iii) there is often no information about what the outcome would have been if another decision had been taken; and (iv) most important decisions are unique and therefore provide little opportunity for learning (see (Einhorn, and Hogarth, 1978)). The conditions for organizational learning are hardly better. Learning surely occurs, for both individuals and organizations, but any claim that a particular error will be eliminated by experience must be supported by demonstrating that the conditions for effective learning are satisfied."

Because of the nature of the feedback structures, often people do not receive timely or appropriate outcome feedback; therefore, learning does not occur.

Rasmussen et al. (1981, pp. 150-51) describe research on human errors suggesting that many errors can be seen as "unsuccessful experiments in an unkind environment." They state:

"Human variability is an important ingredient in adaptation and learning, and the ability to adapt to peculiarities in system performance and optimize interaction is the very reason

> for having people in a system. To optimize performance, to develop smooth and efficient skills, it is very important to have opportunities to "cut corners", to perform trial and error experiments, and in a way human errors can be considered as unsuccessful experiments with unacceptable consequences. Typically they are only classified as human errors because they are performed in an "unkind" work environment. An unkind work environment is then defined by the fact that it is not possible for a person to correct the effects of inappropriate variations in performance before they lead to unacceptable consequences."

Regarding adaptation, Rasmussen states,

> "...in general, the only information available to the person to judge the proper limits of adaptation will be occasional mismatches of behavior and environment. In this way, conscious as well as subconscious experiments are part of the adaptation mechanisms at all levels of cognitive control."

Hence, on all levels of behavior -- skilled, rule-based, and knowledge-based -- people will attempt to ascertain how systems will behave under varying conditions. In fact, without this information, people cannot know fully how to control the system because their mental models will be incomplete.

Dynamic situations may provide the opportunity for learning and error correction, but they do not guarantee it. Decision appropriate feedback, outcome feed-in, may be delayed or degraded such that learning will not occur or may be inaccurate. Therefore, in order to understand human decision making it is necessary to recognize that human adaptability and learning affect the basic feedback control loop and that system characteristics influence the human's ability to learn about systems.

Considering the importance of the sequential nature of decision making, it is surprising that only a few behavioral decision theorists have recently begun to study judgment and choice as a *continuous process*. As noted earlier, the majority of theorists have focused on decisions as *discrete incidents*. Few researchers have questioned whether their findings can be applied to continuous processes (Hogarth, 1981). Those researchers fall into two groups: the pioneers working in the 60's and early 70's and the contemporaries aided by readily accessible and more powerful computer technology giving the field a new surge in the 80's.

The Pioneers

Rapoport (1972, and 1975) provide an overview of the early literature. Edwards (1962) and Toda (1962) are widely credited as the pioneers in this tradition, and Rapoport (1966a) and Ebert (1972) represent the major experimental contributors to the field. The early dynamic decision theorists viewed dynamic decision making through the framework of multi-stage control problems.

Edwards (1962, p. 60) describes the framework of multi-stage control problems:

> "...decision makers are conceived of as making sequences of decisions. Earlier decisions, in general, produce both payoffs and information; the information may or may not be relevant to the improvement of later decisions. The objective of the decision maker may be taken to be maximization of total profit over the long run. But it is quite likely to be desirable to give up short-run profit in order to increase long-run profit. The most common

> instance of such a conflict would arise in situations where some courses of action lead to more information and less profit, while others lead to less information and more profit."

Rapoport (1975) relates Nemhauser's (1966), more specific description of multi-stage control systems as those that can be characterized as several single stage decision systems linked together in a series in which output of previous stages may affect subsequent stages. A single stage decision system is a time slice of a system that can contain six factors:

1) an input state
2) a possible external event
3) a decision variable
4) an output state
5) a stage return (or payoff)
6) a hidden sixth factor, the stage transformation

The decision maker chooses a course of action, the decision variable, and receives some payoff after the stage transformation occurs. In the dynamic situation, previous stages provide the input state to the new stage. Therefore, previous decisions affect subsequent stages. Rapoport (1975, p. 349), concurring with Edwards (1962), summarizes the paradigm as follows:

> "...decisions are made sequentially in time; the task specifications may change over time, either probabalistically or deterministically, and either independently or as a result of previous decisions; information available for later decisions may be contingent upon the outcomes of earlier decisions; and implication of any decision may reach into the future."

The pioneers, (Ebert, 1972; Rapoport, 1966a, 1966b, and 1967) all share a common view of the decision maker as a "constraint optimizer." Through various manipulations to the task environment, these investigators seek to describe differences between optimal performance and actual performance, attempting to ascertain the boundaries of the decision maker's capabilities. This approach is in sharp contrast to that of more recent researchers who seek to understand the heuristics and biases of the decision maker (Hogarth, 1987). The early approach, therefore, sought to describe human performance by comparison to optimal solutions. The tradition was based on subjective expected utility theory which deemed the decision maker as one who always attempts to maximize some gain.

The factors investigated in these studies included:

1) levels of uncertainty (stochastic input vs. deterministic input) (Ebert, 1972; Rapoport, 1966a)

2) effects of varying horizons or stages that the models utilized in calculating their next choice of inputs Rapoport (1966a, and 1966B)

3) whether or not performance can be described by a Bayesian type decision rule (Rapoport, 1966b)

4) realism (use of a more realistic task, i.e. a stock-adjustment task), thereby increasing external validity (Rapoport, 1967)

5) effects of levels of information about the distribution of stochastic external inputs (Rapoport, 1967)

6) effects of stochastic variables in the presence of multiple control choices (Ebert, 1972)

The early research in dynamic decision making focused primarily upon the outcome feed-in link as they deal with the effects of uncertainty of information within that link (Factors 1,5,6 above).

Of particular interest are the studies of Rapoport (1966a, & b). He examines how people perform control on an unstable process of the kind $x(t+1) = a*x(t)$ where $a>0$. Rapoport compares the decision maker's performance with a normative model developed from dynamic programming techniques and later with another normative model, the Bayesian adaptive model. Both studies investigate the effects of uncertainty.

In both studies, subjects are asked to pretend that they are businessmen responsible for purchasing and selling a number of shares of stock per month for their corporation. Costs accrue for decisions made and stocks that remain on the market after decisions are made; in other words, costs accrue for control actions and deviations of x from its set-point. People are told to minimize total costs.

Rapoport (1966a) investigates the effects of uncertainty: a deterministic case versus a stochastic case. The deterministic case implies that all of the variables are specified to the subject and there is no uncertainty about external influences on the system. The stochastic case involves a random external input to the system in the form of demand; therefore, the subject is unable to predict the subsequent state of the system. Rapoport (1975, p. 359) explains,

> "...that the effectiveness of subjects' decision behavior was reduced by half when a stochastic variable was introduced into control task."

Analysis reveals that people do not perform optimally under the model's assumption that the subjects consider all of the stages the system will go through; therefore, he changes the constraints on the model, namely the number of stages taken into account when determining the number of shares of stock to buy or sell. With the change to the model, he finds that a 3 stage horizon fits a small portion of the data: 3 subjects of seventeen. He speculates, however, that the horizon of 3 would vary across decision maker's and would increase to an upper bound with increased experience.

Rapoport (1966b) explores the descriptive validity of another normative model: an adaptive, Bayesian model. The adaptive model assumes that subjects utilize a Bayesian approach to learning the probability distribution of the random external inputs to the system. In addition, he assumes that the decision maker does not consider all of the stages when formulating a particular decision; Rapoport (1966a) reveals that subjects were probably utilizing a low horizon in determining their control choices. Therefore, the model is built with an unknown stage duration in which to calculate optimal solutions. He finds that subjects' median decision behavior can be described well with the Bayesian (adaptive) optimal model. Analysis of the subject's data reveal that a horizon of j=7 stages provides a good fit to the data, but that varying values of j are needed to model people's behavior. He concludes,

> "It seems that human multistage decision making behavior may be characterized as constrained optimal, provided the structure of the task is fully understood and provided the experimenter knows the perceptual and intellectual constraints operating in the situation and imposed upon decision makers (Shuford, 1964). Within the realm of research on decision making such constraints will be concerned with short-term and long-term memory, information processing, overestimation and underestimation of probabilities, future implications of present decisions, etc." (Rapoport, 1966b, p.60)

He adds that the then current knowledge of constraints was limited to average constraints. While the model used has the advantage of being analytically tractable, this type of unstable process, $x(t+1)= a*x(t)$ where $a>0$, has no real-world counterpart.

In a later study, Rapoport (1967, p. 195) thus investigates inventory-control problems and thereby extends the previous studies by,

> "...treating a class of decision problems which are more meaningful to DM [decision maker]; because they attempt to mirror the key features of certain concrete situations, they are presumably of greater interest to the behavioral scientist."

This class of problems represents what system dynamicists know as stock-adjustment problems (Sterman, 1989a, and 1989b). These problems involve a decision maker controlling a level of stock such as inventory, water, or virtually any quantity, in the face of either a constant or varying demand. Costs may accrue for production changes, storage, and shortages. Rapoport (1967) hypothesizes two models of ordering policies under two treatment levels of information:

> hypothesis 1) an ordering policy that minimizes total expected costs over the future
>
> hypothesis 2) an ordering policy that minimizes immediate expected costs
>
> treatment 1) complete information about the distribution of demand values
>
> treatment 2) no information about demand

Subjects' performance is invariant with respect to the two treatment levels of information. In addition, he finds that neither ordering policy is supported by the results, and the data suggest that subjects ordering policies are not constant over trials; stock orders slightly decrease over trials and are strongly related to the previously observed demand, despite the fact that subjects are told that demand is independent on each stage.

Because previous research had indicated inconclusive results regarding the effects of stochastic variables in multi-stage problems, Ebert (1972) looks at the effects of uncertainty and multiple control choices in the "aggregate scheduling problem", a stock-management task with multiple control choices. Previous work had not investigated multiple control choices. The task is similar to the above stock-adjustment task, but it employs two control inputs: workforce and production rate. Subjects' goal is to minimize total costs in a dynamic situation in the face of varying demand. Ebert (1972, p. 237) found:

> "...the results indicated that the presence of uncertainty in the two-variable, control task does not impair decision performance to the extent that has been suggested in one-variable decision studies."

This contradicts the results of previous studies that noted a 50% decrement in performance at the introduction of a stochastic variable (Rapoport, 1966a). Since Ebert's study is focused upon the uncertainty issue, no specific hypotheses are made concerning the multiple decision variables themselves. However, more variables may imply greater system complexity, and Ebert relates that Peterson and Beach (1967) showed that more complex tasks may afford increased performance as compared to simpler tasks.

Several observations can be made from the previous empirical studies:

1) findings concerning uncertainty are inconclusive

2) people utilize varying horizons when formulating their strategies

3) viewing the decision subject as a "constraint optimizer" has proved of limited use with unclear results

4) people's strategies in stock-adjustment tasks varies from the behavior as prescribed by optimal ordering policies

5) information about the demand distribution does not always increase performance in a stock-adjustment task.

It remains unclear as to what strategies people are actually using. The insistence of early work on viewing the subject as a "constraint optimizer" leaves little room for explaining the processes of decision behavior. Rapoport (1975, p. 361) himself expresses his doubts but tries to rescue the approach by stating:

> "The research methodology proposed above calls for employing the optimal solution as a base line, then modifying it by placing psychologically interpretable constraints on DM's [decision maker's] perceptual or cognitive processes. The time horizon discussed above is one constraint, but there are many others. But even before attempting to constrain the optimal model the investigator should realize that there are as many optimal models as there are objective criteria. Moreover, even if the objective criterion is explicitly stated by the experimenter there is no assurance that a normal subject will understand it or attempt to maximize it if he understands it, nor is there a simple procedure for determining which criterion or criteria the subject adopts or whether or not he shifts criteria during the task....It seems advisable, therefore, to assume various reasonable objective criteria, derive, if possible, the optimal decision policy under each of these assumptions, and whether or not the optimal policies are further constrained, test DM's behavior against each of the resulting models."

In addition to these criticisms, one may question the generality of the results. It is not easy to relate the findings at this point in the tradition to how people perform decision making in everyday settings.

Although interest in dynamic decision theory has continued from the 60's through the end of the 70's (Broadbent, and Aston, 1978; Mackinnon et al., 1980; Rapoport, 1975), dynamic decision theory has not been a very active research area until lately. This quiet period in the research could be attributed to the above mentioned limitations and to five other possible causes outlined by Rapoport (1975, and 1972):

> 1) The mathematical sophistication of dynamic decision problems and the need for time consuming computer programs have contributed to the lack of research in the area.
>
> 2) "Analyses of human and animal behavior have been mostly predicated on the assumption of a one way dependence of responses upon stimuli, ignoring the equally important dependence of stimuli upon responses according to the organization of the environment in which responses are emitted. But without explicitly considering this latter dependence dynamic decision problems cannot be solved and goal-directed behavior cannot be properly explained."
>
> 3) The tools of feedback theory and control theory were not well known to behavioral scientists; therefore, they could not explore the dependence of stimuli upon responses.
>
> 4) Designing and implementing experiments in which the task characteristics change as a function of the subject's decisions without the availability of a small on-line computer are difficult to construct.
>
> 5) "...Characterizing the concepts of dynamic decision theory, control theory, or system theory, so as to successfully delineate segments of the environment plus a sequence of decisions as a dynamic decision task," is difficult.

The Contemporaries

Partly in response to the ready availability of computer technology there has been a new surge in the field in the 1980's. While addressing some of the limitations faced by the forerunners in the area, psychologists and behavioral decision theorists have recently begun to re-examine dynamic decision making.

Dörner (1975) pioneered the investigation of human problem-solving and decision making in complex, dynamic computer simulations with his original simulation, TANALAND. His work represented the first of the attempts to look at human performance in more "real life" circumstances. His initial studies led to a wave of explorations into the influence of various personality characteristics, situational factors, and the like on performance in computer simulations.

Dörner himself looked at the extent to which different levels of IQ influenced problem solving behavior. TANALAND simulates an African landscape including two groups of people and contains over 50 variables which are interconnected by a series of positive and negative feedback loops. Subjects' goal was to improve the living conditions of the population through various agricultural decisions. It was found that subjects' cognitive ability to cope with complex systems was inadequate. The main deficits Dörner found in the subjects' behavior were their inability to deal adequately with exponential growth and the interdependencies between variables. He furthermore found that subjects tended to focus narrowly on just a few variables and did not wish to deviate from their set course of action. As a consequence, subjects destroyed the originally stable eco-system. Dörner describes as due to "linear thinking" and suggested more emphasis on thinking in causal networks.

Dörner, Kreuzig, Reither, and Stäudel's (1983) work culminated in LOHHAUSEN, which contained over 2000 variables, multiple positive and negative feedback loops, and substantial delays. In these efforts, they explored how people, playing the role of mayor, perform in controlling this simulation of a small town. Subjects had difficulty balancing short-term goals with long range goals. They found a lack of cognitive perseverance, i.e. high performers test one causal hypothesis completely, while low performers tend to jump from one hypothesis to another. Dörner called this jumping phenomenon "thematic vagabonding," or generating new hypotheses without testing them appropriately. In the presence of delays, much time can pass between control actions and system responses. Disconfirmation of incorrect hypotheses can only be achieved by waiting for the feedback of those control actions; most people do not incorporate these delays into their search for appropriate control policies.

Following these pioneering studies and continuing in the psychologist's tradition, Reichart (1986), who was interested in the effects of personality characteristics on the ability to control complex processes, devised COLD STORAGE DEPOT. Subjects were asked to control the temperature of a system with a steering wheel. Subjects had difficulty with the time-delay of the nonlinear function driving the steering controls; some subjects planned for it, while others only changed their strategy immediately after feedback.

DAGU, devised by Reither (1981), consisted of a climate and population simulation of an African area who required aid. In this study, novices were compared to experts that were trained in how to administer aid in these types of situations. It was found that novices thought more in causal chains than in causal webs, meaning that they did not take side effects into account, but only thought of main effects. Experts out-performed the novices by considering the consequences of the actions beforehand and applying their strategies consistently. However, even the experts were not able to stabilize the variable "population size," one prime example of an exponential growth process (Funke, 1988).

Funke (1988) reviewed these and similar studies motivated by Dörner's work; the basic findings of much of his review can be summarized as follows:

1) subjects have problems predicting and controlling variables that grow exponentially

2) subjects exhibit problems incorporating delays into predictions of outcome of control actions

3) subjects have problems with complex connectivity of variables

4) subjects remain focused on only salient variables

5) subjects do not recognize the consequences of their actions.

Kluwe, Misiak, and Haider (1989) explore the processes of information utilization in dynamic decision making thereby bridging the gap between psychologists interested in personality characteristics and researchers focusing on the effects of different task characteristics.

They postulate three stages of learning: 1) information orientation, 2) exploration of alternatives, and 3) system control. Information orientation is concerned with learning what variables are present and initial classification of information. Exploration of alternatives is an extension of the first stage with more detailed exploration of the ways variables interact or influence one another and the effects of control inputs. The third stage, system control, is characterized by the subject's ability to integrate many system variables and generate effective strategies for controlling unstable processes. Heuristics in the three stages are distinct.

Parallel to the renewed interest among psychologists, behavioral decision theory experienced a similar surge in interest for dynamic decision making. While the early picture of the decision maker was as a constraint optimizer, Tversky and Kahneman's (1974) new paradigm, shifting from a normative field towards a descriptive field, began viewing the decision maker as a limited information processor using heuristics.

Motivated by Hogarth's (1981) appeal to apply Tversky, et al.'s concepts to dynamic decision making, Kleinmuntz (1985) tested the performance of models of cognitive heuristics including a production system under varying task characteristics. He used a medical diagnosis and treatment task in which subjects test, treat, and try to cure patients. Some of the results indicate that informational aspects of the task, (symptom diagnosticity, disease base-rates), do not influence performance of the heuristics as much as feed-back characteristics of the task. He concludes:

> "The ability to select relevant variables seems to be more important than procedural sophistication in the processing of that information." (Kleinmuntz, 1985, p. 696)

Hence, the important ingredient is the "conceptualization" of the task rather than the computational effectiveness of the subject.

In a later study using the same task Kleinmuntz, et al. (1987, p. 342) investigated the cognitive effort/accuracy tradeoff: the ways in which the task demands influence the choice of decision strategies. They note:

> "...in many situations, a decision maker can dispense with judgment and simply act. Not only is the cognitive effort that would normally be expended upon judgment conserved, but, because of the opportunity to catch and correct poor choices, the outcome is acceptable."

Viewed from the closed-loop framework proposed in this thesis, Kleinmuntz's studies shed light on the action feed-out link. As Hogarth (1981) had earlier explained, the level of *commitment* to any one choice is reduced because there will be opportunities for the person to correct the choice later. Results of this study indicated that people were typically relying upon judgement-oriented strategies instead of action-oriented strategies.

Judgement-oriented strategies were indicated by the extensive use of diagnostic tests prior to actions. Action-oriented strategies, conversely, imply little to no testing before action is taken, and these were not extensively used even though they would have produced better performance in some situations. There are several possible explanations for their results, so one cannot consider this conclusive evidence for any particular explanation for the strategies used by subjects. However, one possible explanation of the results lies in the use of the "doctor script" which dictates the use of testing before treating: subjects might show a tendency to act according to what they think is "appropriate" physician practice in a task for which they are unfamiliar.

Huber (1986, p. 68) views both the system to be controlled and the controller as open systems and advocates the need to combine the two into one overall system. He, thus, proposes a systematic treatment of both the task and the decision maker. Huber relates his motivations as follows:

> "The problem with the decision theoretical approach in research on multistage decision problems is that it does not reveal information about the process of decision making and about the interaction between the decision maker and the components of the task. This is a very unsatisfactory situation because in other areas of decision theory process models are getting more and more important."

Specifically, Huber investigates the extent to which mental representations, which may be described in terms of completeness and adequateness, and other factors, such as "locus of control," influence dynamic decision making. He explores the role of the mental representation and personality characteristics in a multi-stage betting game.

A multi-stage betting game is analytically similar to the multi-stage decision theory paradigm, with the exception that the decision maker begins with some amount of capital and gambles a certain portion, a stake, in each stage under varying task characteristics, such as probability of winning or losing.

Using an ecological cover story and a treatment variable of increasing probability of winning, Huber found that subjects characterized as having an "internal locus of control" perform better than subjects characterized by an "external locus of control." Locus of control gives some indication of the belief subjects have in their control over system variables. In addition, he found evidence that the mental representations of the subjects are closely related to the generation of goals for controlling target variables within the system.

Brehmer (1989, and 1991) investigated how people control self-reinforcing processes with the introduction of delays in control and outcome feedback. Delays can arise from several sources in the feedback loop: from control action to implementation and from implementation to the final outcome. In a simulated fire-fighting task, people were asked to allocate fire-fighting units to simulated forest fires that expanded exponentially. Delays were introduced in the control structure and in the reporting of fire-fighting units' progress on the fires. In order to perform the task efficiently, Brehmer states that it is important for people to adopt a feed-forward strategy in controlling the system. In other words, people must incorporate the delays into their mental representation and appropriately utilize observed feedback, which lags behind the actual or current state of the system.

Another factor investigated in Brehmer's study was the effect of observable delays versus hidden delays. Observable delays are those that can be seen as a natural consequence of the system; people expect some delay in the mobilization of fire-fighting units. Hidden delays, those in the reporting of the progress of fire-fighting units, are not necessarily expected; people expect that a unit's position and progress will be reported in a timely fashion. The results suggested that people had the most difficulty with delays of the second type: hidden. Their performance suggested that they did not compensate for the delayed reporting of the fire-fighting units; they treated the displayed information as if it were indicative of the current state of the system. Thus, the results imply that "all delays are not created equal" (Brehmer, verbal communication with the author).

Sterman (1989a, and 1989b), consistent with the newer dynamic decision paradigm exemplified in Kleinmuntz (1985) and Brehmer (1989) brings into sharp focus the importance of the feedback structure of the task. Sterman's overriding concern rests in how people perceive the consequences of their own actions in a feedback context. He finds that people typically misjudge the influence of their own actions. Take, for instance, supply-line, phenomena caused by the existence of delays. If a retailer places an order for some commodity, the supplier takes some time to fill the order and to ship it; therefore, there is a delay built into the ordering system. In experimental settings, people typically tend to ignore this information, resulting in instability. Sterman provides an operational theory of feedback that can readily be tested.

He performed two experiments that are similar and produce similar conclusions. Sterman (1989a) investigated a simulation of an aggregate capital-producing sector of the economy in which one player is asked to manage a stock-adjustment task. Sterman (1989b) explored a one-sector goods producing economy in which multiple players act as teams to meet a one-time increase in demand. The findings can be summarized as follows:

> 1) People tend to ignore the supply-line; subjects don't sufficiently account for orders already placed with their suppliers. Large oscillations are the consequence. Once the double ordered goods arrive, subjects are faced with large inventories and need to cut their orders well below the equilibrium rate.
>
> 2) People have difficulty controlling exponential growth or self-reinforcing loops: they don't recognize the fact that they must respond rapidly and massively to a demand surge.
>
> 3) Subjects misjudge the effects of non-linearity; people don't recognize that the possibility for control and the effects of control can vary significantly once the system is perturbed from its "normal" state.
>
> 4) People utilize open-loop mental representation and blame exogenous factors for their performance.

These findings are similar to those of the studies mentioned earlier (Brehmer, 1989; Dörner et al., 1983; Kluwe et al., 1989). Rather than seeing people's dysfunctional behavior as isolated errors, Sterman (1989a) interprets his findings within his framework of "misperceptions of feedback". He views his findings as examples of:

> "...a failure on the part of the decision maker to assess correctly the nature and significance of the causal structure of the

> system, particularly the linkages between their decisions and the environment." p. 18

Sterman encourages a more systematic exploration of the misperceptions and calls for a better understanding of the interactions between subjects' intendedly rational decision rules and the feedback structure of the task to be controlled.

1.3.5. Summary

In different ways, practitioners in each of the four fields discussed have contributed to an improved and more comprehensive understanding of decision makers' effects on, and influences from, their task environments. Supervisory control investigates feedback perceptions and decision making in a specific task environment, that of controlling machinery. Experimental economics examines the effect of action feed-out in the form of trading mechanisms from the marketplace on attaining the goal of market equilibrium. Behavioral decision theory examines decision makers' perceptions of outcome feed-in from their task environments and their ability to interpret and act upon information. Dynamic decision theory helps explain the consequences of decision feedback on the task environment and the decision maker over time. In all of these fields, the need for a taxonomy of tasks that defines and categorizes decision making, the task environment, outcome and actions has been recognized. Very often the struggle for a taxonomy has been phrased as the need to characterize the "complexity" of the task environment. Rapoport concurs with Shuford (1964, p. 60) who concluded,

> "It seems that human multistage decision making behavior may be characterized as constrained optimal, provided the structure of the task is fully understood and provided the experimenter knows the perceptual and intellectual constraints operating in the situation and imposed upon DMs."

However, researchers at that time only had limited knowledge of the constraints of the task environment and the person; consequently, they viewed the decision maker as a "constraint optimizer" which limited the number of alternative explanations of behavior. The paradigm, however, did elucidate the general nature of the decision environment.

Funke (1988), in his review of the psychological simulation studies, describes the ambiguity of the concept of complexity as he groups the simulation experiments into three categories: up to ten variables, up to 100 variables, and over 100 variables. He explains that complexity cannot merely be considered on the basis of the number of variables alone within a simulation. For instance, is a system with 24 variables more complex than a system with only 12? Or a system with over 100 variables 10 times more complex than a system with only 10? He argues that complexity needs a more rigorous definition.

Mackinnon (1980), early advocates of characterizing the task environment more systematically, study the effect of three "complexifying" factors: (i) number of elements in the system, (ii) connections between them and (iii) the presence or absence of random variation. To their surprise, the three "complexifying" factors demonstrate either insignificant differences, or significant differences in the direction opposite to that predicted. They reject the notion that increases in the number of connections necessarily increases complexity by pointing out potentially beneficial effects:

> "Connections between elements of a system, and the indirect connections formed over time, may be a source of stability, or produce correction from subsystem to subsystem under a range of conditions." (Mackinnon et al., 1980, p. 294)

Since more complex systems, as measured by their "complexifying" factors, did sometimes attain better performance than simpler systems, they suggest rethinking the effects of complexity. Bakken (1990) gives a nice example illuminating the discussion of complexity:

> "...imagine an old wood stove and a modern thermostat controlled oven. An oven controlled by a thermostat is more complex than a manually operated one, yet it is easier to control and improves cooking for most novice and medium level chefs. In other words, system complexity will of course improve performance if "complexity" means well designed control process."

Thus system complexity can take many different forms and influence decision makers, their goals, their actions, the task environment, and the resulting outcomes in different ways. In a later study, Mackinnon, et al. (1985, p. 160-1) criticize both the dynamic decision tradition and the psychological-simulation tradition similarly:

> "Thus whilst DDM {dynamic decision making} would appear to be a relevant area of inquiry in a world in which many significant decisions

> are sequential and dependent, research in the area has been confined to a small number of specific tasks, each of unknown generalizability and exploring a limited sector of the universe of task systems."

They describe a preferable approach that:

> 1) Enables and encourages a systematic exploration of the universe of dynamic task systems.
>
> 2) Permits a precise and formal description of the degree of difficulty or complexity of the task.
>
> 3) Allows tasks to be compared.
>
> 4) Is capable of providing "base-line" information of human ability in these tasks.

The dimensions of complexity as defined by the fields that examine decision making can be generalized and categorized according to their effect on the five elements of the decision making system, namely the decision maker, the desired goal, the action feed-out, the task environment, and the outcome feed-in. The issues discussed in the various studies can be categorized by the following dimensions.

Time pressures can directly influence the tasks of the decision maker because all tasks occur within the context of time. For example, the decision maker may not have enough time to both perceive goal discrepancies, to formulate a strategy, and to take an action. Further, the action feed-out link and the outcome feed-in links may contain delays that will make the relationship between decisions and resulting outcomes more complex.

The number of cues or amount of information reaching the decision maker affects complexity. At any given time there may be multiple cues from which the decision maker can infer system behavior. The number of cues available is a function of the quality of the outcome feed-in link. For example, the total number of cues available, the number of cues sampled, the information display of cues, and the amount of time available for sampling and display all influence the quality of the outcome feed-in link.

Discrepancies among cues or information may challenge the decision maker. At any given time, the decision maker must attend to multiple variables that must be maintained in order to achieve system goals. Discrepancies in outcome feed-in cues will force the decision maker to make trade-offs in expected outcomes for some goals in order to more closely approach desired outcomes for other goals. For example, operation of a nuclear power plant requires management of myriad sub-goals which may have discrepancies that must be maintained within appropriate boundaries in order for the overall power plant to achieve optimal capacity-cost-safety benefit trade-offs.

The number of action feed-outs within a system will influence feedback characteristics. Most real systems require several choices to be made in order to achieve stability. For example, businesses managers must decide upon appropriate production rates, personnel levels, marketing expenditures and quality goals in order to achieve a desired goal of increased cumulative profits.

Uncertainty can affect both the outcome feed-in and action feed-out links. Usually, not all outcome feed-in from the task environment will be caused by the decision maker. However, the decision maker may be unable to separate causal outcome feed-in from randomness. The task environment can be influenced by elements other than the decision maker and not all outcome feed-in to the decision maker is produced by the task environment. The outcome feed-in link may contain randomness that would require a systematic bias in output. Decision makers must realize that there are multiple agents interacting with the system. In order to produce high quality of overall system performance, each agent must take into account the actions and effects produced by other agents. Furthermore, the way that the individual agents exert control and the way that action feed-in from a given agent becomes outcome feed-out for another agent determines effective organizational structure.

Incentives can influence the desired outcome, relative importance of some goal discrepancies over others, preference for particular actions, and perceptions of particular environmental cues. For instance, a politician may give more weight to issues that affect a major contributor to his campaign, and he would in turn be more cognizant of literature pertaining to those issues and may be more likely to vote in the major contributor's interests.

The structure of the organizational context of dynamic decision making may be important. Organizational structure may constrain or impede the decision maker's actions. The freedom of relating outcomes to actions, i.e., the number of rules that are imposed on the decision maker, can influence the decision feedback loop greatly.

The training and experience of decision makers teaches them how to perceive outcomes, to understand consequences of actions, to understand how the environment transforms actions to outcomes, and to know how to best map actions into outcomes. Training and experience can influence the decision maker's ability to make decisions that approach his desired outcome.

Numerous other dimensions of complexity exist that have received less attention in the literature; they include: number of states present, degree of non-linearity, system stability, and eigenvalues and eigenvectors of the system. In addition to the dimensions that are a part of the complexity of the physical system, people's information acquisition and use add additional dimensions of complexity along the lines of saliency, observability, and measurability.

Sterman (1989a, and 1989b) suggests that the consistent findings of the "misperceptions of feedback," for instance, the misperception of delays, exponential growth processes, and non-linearity, provide a beginning for a classification system of dynamic decision task characteristics. A comprehensive approach, for other researchers, would ascertain the effects of other dimensions of the task environment and determine their precise effects upon performance in systems of varying degrees of complexity. Currently, comparison of results of dynamic decision tasks across the various research domains are difficult since many researchers select their task variables without wide regard for how their study will aid the development of a theory. Researchers, additionally, utilize a task system that is operationalized to suit their unique purposes. Therefore, our understanding of dynamic decision making is limited by our unsystematic treatment of the task environment.

For obvious reasons, we cannot investigate all of these dimensions of complexity in this thesis, though all of the above dimensions could be fruitfully investigated within the essential control loop framework. One of the dimensions of complexity our research will examine is the effect of delay on achieving goals. Many of the researchers cited include some dimensions of delay. The effects of delay have been documented in many of the previous studies; however, no studies have systematically varied the effect of varying delay lengths and shown the specific influences of delay on the self-regulating feedback loop.

Another very important but widely ignored dimension of complexity is the effect of side-effect feedback in the task environment. Gain determines the strength of unintended outcome feed-in resulting from the decision maker's action feed-out. Decisions may be made that influence aspects of the task environment other than those intended resulting in unexpected outcome feed-in to the decision maker.

Further exploration of these two dimensions provides an opportunity to test the effectiveness of the framework through verification and clarification of previous findings and an opportunity to provide baseline data of human performance in the framework of complexity proposed.

2. Experimental Design

2.1 Design considerations

The investigation of people's information use may take two forms: examining decision rules under relatively simple task demands and examining decision rules under very complex task demands. Dörner, Kreuzig, Reither, and Stäudel's (1983) work represents the latter approach that has shown some success in describing decision biases under various complex settings. Unfortunately, relating their findings to novel situations can be done in only a rough sense, since we do not have a means to effectively compare situations of varying complexity.

Considering the vast variety of physical systems in existence, it is a daunting task to develop a general theory of people's information use. We may find comfort, however, in the findings reported in the system dynamics literature that there are certain generic structures that seem to govern a vast majority of the decision tasks with which subjects are confronted. It is my presupposition that once we understand how decision makers go about their decision tasks in "simple" systems, we have laid the foundation that we can expand upon to understand decision makers in systems that consist of a combination of one or more of those structures. One generic structure that is of particular importance is the broad class of stock-adjustment tasks.

2.2 The importance of stock-adjustment tasks

A vast number of human activities can be characterized as attempts to control a stock and maintain its value close to a target value (Sterman, 1989b). Stock adjustment problems are prevalent on different levels of aggregation. Individuals change their car's velocity to drive at a desired speed, regulate the water's temperature to shower comfortably, and try to eat in such a way as to maintain a desired weight. Companies borrow money to manage their cash balance, increase production to keep enough goods in inventory, and hire employees to meet their labor needs.

In all of these examples, the objective is to maintain a stock at its targeted value in the presence of disturbances such as losses, usage and decay. Indeed, stock-adjustment, goal seeking is so prevalent that people attempt to control using this principle even when it is inappropriate to do so. Consider an example from health psychology:

> "One straightforward problem stems from the fact that certain physical disorders (e.g., hypertension) have no known symptoms. Despite this, people with the disorder persist in attempting to ascertain their present condition by means of easily monitored (internal) perceptual events. In effect, they define for themselves some symptoms to observe. Of greater importance, they then proceed to regulate their behavior on the basis of the presence or absence of that symptom (Leventhal, Meyer, and Nerenz, 1980).
>
> This tendency leads to either of two kinds of problems. If the symptom goes away, people may stop taking their medication, believing that their blood pressure is down when it is not. If the symptom persists, on the other hand, people may become despondent, believing (incorrectly) that their blood

> pressure is remaining elevated despite their faithful following of the doctor's orders." (Carver, and Scheier, 1982, p.127 author's parentheses)

Since stock-adjustment problems are so widespread and virtually ingrained in people's heuristics, they represent a good starting point for developing a data-base of performance in dynamic systems.

2.3 The task

Having motivated the importance of stock-adjustment tasks and delay and gain as treatment conditions, we now describe the particular instantiation we have chosen for our study.

The cover story chosen for the stock-adjustment task is that subjects are to manage an inventory production system, in which the stock to be controlled is inventory. The following instructions are given to subjects:

INSTRUCTIONS

1. Objective

You are the inventory manager of a company. Your responsibility is to minimize the total costs your department will accrue during the game described in this set of instructions. The inventory you have to manage is decreased by sales and increased by production. You know only partially what your sales volume will be but you have full control over your production. Coordinating production with sales while avoiding inventory gaps is the central problem you face in this game.

2. Cost structure

Costs arise from two sources: production changes and inventory gaps. To minimize the total costs, you need to know the exact nature of these costs:

- You try to maintain an inventory of 0 units. If the actual inventory is larger, too many units are held in stock and cause unnecessary inventory holding costs. If the actual inventory is smaller, goods are not in stock when the customer demands them. This causes irritation for the customer and losses for the company. In the game, it is assumed that these costs are quadratic: Large changes cost considerably more than small changes.

- Production changes in a real company imply such costs as hiring, training, or firing people. Again, it is assumed that the costs of production changes are quadratic. A change in production is twice as expensive as an inventory gap.

3. Information on sales

Unanticipated changes in sales will cause an unwanted positive or negative inventory gap. You are facing two kinds of sales: sales independent of your actions and sales dependent on your own actions.

Independent sales follow a random path. **Your best bet is to expect that independent sales in the next period will be the same as they are in this period.** Actual sales in the next period, however, may differ anywhere from plus 15 units to minus 15 units from what they are in the current period.

Dependent sales are influenced by your own production decision. Production can influence sales in two opposite ways:

- As you increase production, you implicitly increase your workforce and pay a higher wage sum. A higher wage sum increases demand and leads to higher sales for yourself.
- As you increase production, competitors become aware of the increased opportunities and increase their own sales efforts: Increased competitor activity causes some of your customers to cancel some of their orders that they had placed with you.

In any particular game, you will only encounter one of the two sales effects.

4. How you participate in the game

A game consists of 30-40 decision periods. In each period you have to decide by how much you want to increase or decrease production. The time it takes for a change in production to take effect can vary from game to game. In the course of the 4 sessions, you will play 15 different games altogether.

Figure 2.1. illustrates the task and the two treatments: delay and gain.

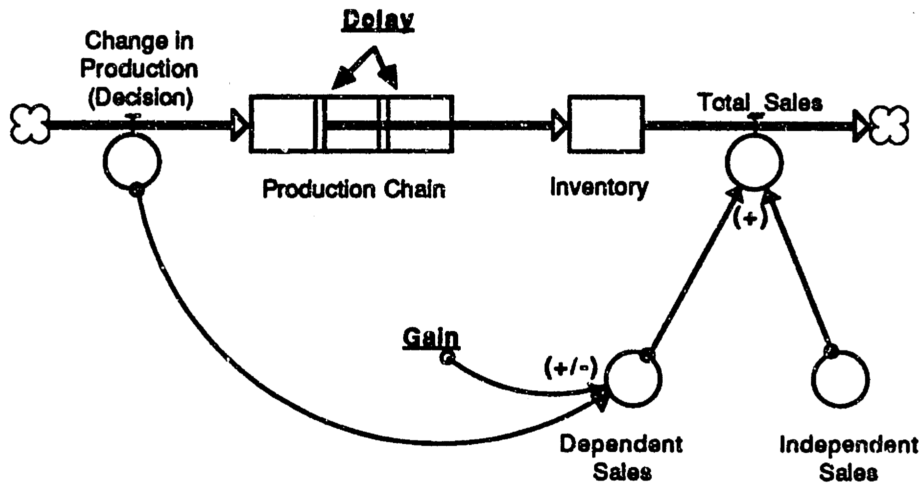

Figure 2.1. Stock-adjustment task with two dimensions of complexity: delay and gain

At the center of Figure 2.1. is a rectangle which represents inventory. Inventory may be an accumulation of any quantity. The pervasive accumulator concept is central to the stock-adjustment task and may take various forms, such as any structure that accumulates in-flows and out-flows. For instance, body temperature is an accumulation of all of the temperature loss and energy gained within the human body, or a bathtub is an accumulation of the water in-flowed and out-flowed. In the surface structure we have chosen, the accumulation to be controlled is inventory -- the accumulation of goods that have been produced minus all of the goods that are sold.

Examining the in-flow to inventory on the left of Figure 2.1., we see a production chain. The user is in charge of determining the total production amount through changes in production. Until produced goods are finally added to the inventory, they may go through different stages; those different stages are actually one of our treatment conditions: delay. Once we commit ourselves to taking delivery from a supplier, we initiated an event which will change our inventory even though the inventory has not yet physically increased. If we lived in Europe, we would receive our oil in perhaps weeks, while if we lived in the United States, we would receive our oil later, in perhaps a month or two. The whole process taken together we would call a production chain or supply-line. By contrast, if we happened to live in the Middle East we would suppose that our shipment of oil could be delivered virtually within the same day.

The outflow of inventory is sales, located on the right side of Figure 2.1. We have split sales into two components: independent and dependent sales. Sales may be independent of the actions of the decision maker, and may be dependent upon the actions of the decision maker. What is typical of a system is that a control action might have unintended side effects in addition to the intended control action. For example, the more goods you produce in an economy the more salaries and wages you pay out, and in turn, the additional salaries and wages may be used to purchase your own goods. Thus as you increase production, you increase your own sales, an effect on which the Keynesian multiplier system depends. Similarly, one could imagine that as you increase your production you give signals to your competitor of lucrative sales opportunities. Your competitors, at the same time, increase their marketing activities and draw some of the sales away from you. Gain, then, is the sign and strength of those unintended side effects.

Determining the exogenous input to the task requires particular care. Many previous experiments, particularly the ones coming from the control engineering perspective, assume a one-time shock impinging on the system. It is the subject's task to react to the shock and to bring the system back into equilibrium. These tasks can be divided into two experimental conditions: 1. The initial state of the system is in disequilibrium without any further exogenous inputs impinging on the system. 2. The system is in an initial equilibrium and disturbed by a one-time pulse or step change. Both designs frame the decision task to be solved in a special way. The subject is focussed on an end state of the system and the problem is to figure out how to attain the end state as quickly and as smoothly as possible.

The final state focus is well suited for emergency and rescue control tasks. Diagnosing a sick patient and saving her life with adequate surgery would be one example. Bringing a spreading forest fire or an overheating nuclear reactor under control are other examples.

Most stock-adjustment tasks are characterized by an on-going process of control in the face of fluctuating disturbances. Disturbances in the inventory-control problem studied in this thesis are portrayed as independent sales and modeled as random-walk exogenous input. I fully inform subjects as to the nature of the disturbance on the system before the experiment begins and tell them: " your best guess for independent sales is the previous week's independent sales."

To prevent memorizing particular random-walk patterns, 15 different random walk patterns were generated and used. Every subject received the same sequence of random-walk patterns. The objective of preventing memorizing patterns was judged to be of sufficient value to justify one negative side effect of the design decision: Associating a particular random-walk pattern to a trial position leads to a confounding of the practice effect with a possible sequence effect generated by the random-walk patterns.

Every self-regulating feedback control loop operates within a set of objectives. We framed the task as a cost-minimizing task and associated cost with discrepancies of inventory from its set-point and with the amount of control exerted to regulate inventory. The level of deviation of the system from its given set-point is of considerable importance. For example, if body temperature were raised by only a few degrees to 100° Fahrenheit for a short period, the result would not be too detrimental; by contrast, if body temperature were raised to 102° Fahrenheit, the result would be far more damaging. Temperature deviations for different systems from their set-point vary in cost to humans. For instance, it would not be conducive to life if the global temperature of the Earth were raised more than a few degrees. Similar to the varying degree of significance that temperature plays in different systems, our task system implements non-linear cost for deviations from equilibrium. For analytical convenience and to parallel the existing literature, we have chosen a quadratic cost function. Subjects' task is to choose production to minimize:

$$\text{Min} \sum_{t=0}^{t\ end} (a*\text{Inventory}^2 + b*\ \Delta\text{Production}^2)$$

$(a=1;\ b=2)$

2.4. Design and Procedure

It was to be expected that the general competence of the subjects to solve the task might vary considerably, favoring a within-subject design rather than a between-subject design. Since feedback strength and delay length might very well interact, a two-factorial experimental design seemed to best serve the purpose of the study. A Latin-Square design was chosen. We designed the experiment in a way that each subject would receive all the treatment conditions. Rather than to implement a mere delay / no delay treatment we wanted to be able to distinguish between moderate and severe delay. Similarly we wanted to investigate the effects of moderate /severe positive gain and moderate/severe negative gain. A 3(delay)*5(gain) treatment matrix resulted. We implemented delay length of 0,2, and 4 rounds and gain of -0.6,-0.3,0,+0.3, and +0.6. Figure 2.2. illustrates how the two treatments 'Delay Length' and 'Gain' were mapped into one condition and how the condition was in turn allocated to the subjects and trials. Note that each row and each column contains one and only one condition instance.

Subject number	Trial 1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	I	B	N	G	A	F	C	M	H	J	D	L	O	K	E
2	K	D	G	F	C	E	O	J	A	M	B	H	I	N	L
3	L	A	M	O	D	J	G	B	E	N	F	K	C	I	H
4	B	F	J	H	N	D	A	K	I	E	L	M	G	C	O
5	N	O	A	I	M	H	J	D	C	L	G	B	F	E	K
6	C	H	E	M	B	N	K	A	F	O	J	I	L	G	D
7	O	M	I	C	J	L	H	G	D	F	K	N	E	B	A
8	J	E	C	D	L	O	I	H	M	K	N	G	A	F	B
9	F	L	K	J	E	B	N	C	O	H	I	A	M	D	G
10	M	I	H	N	G	K	L	O	B	C	E	J	D	A	F
11	G	C	B	E	O	M	D	L	N	A	H	F	K	J	I
12	D	G	L	K	F	C	E	I	J	B	A	O	N	H	M
13	A	N	D	L	H	I	M	F	K	G	C	E	B	O	J
14	E	J	F	A	K	G	B	N	L	I	O	D	H	M	C
15	H	K	O	B	I	A	F	E	G	D	M	C	J	L	N

Trial number

Subject per sequence table

Delay \ Gain	-0.6	-0.3	0	+0.3	+0.6
0	A	B	C	D	E
2	F	G	H	I	J
4	K	L	M	N	O

Latin Square design matrix

Figure 2.2. Subject per sequence table and Latin-square design matrix for the combination of 3 levels of delay and 5 levels of gain.

2.5. Subjects and Incentives

Seventeen MIT students, fourteen undergraduates and three graduate students, enrolled to participate in the study. Two subjects, for undetermined reasons, did not complete the study: one terminated after the first session and the other terminated after the second session. Their data were examined and no extreme differences from other subjects' data were apparent. Participation required, on average, six hours spread over four sessions in a two to three week period.

Subjects were paid for their participation in the study. Subjects' total pay consisted of a fixed amount of $20 and a performance-based amount. Actual performance-based pay varied between $5 and $45. On average, subjects received a total of $40.

2.6. Protocol

A stock-adjustment system, in the context of a business game and written by the author, was presented on an Apple Macintosh computer.

Subjects were told that they would play the role of an inventory manager of a company. Their task was to minimize total cost of their department, which would accrue during the game. Subjects were told that the inventory to be managed would decrease by sales and increase by production. They were informed that sales could only be partially controlled, since some of the total sales were determined randomly, while they had complete control over production. They were asked to coordinate sales and production, while avoiding inventory gaps (see instruction sheet earlier in this chapter).

Subjects were given complete information about the system. This meant that subjects were informed as to the stochastic nature of independent sales and how all variables were related and calculated.

Week	59	60	61	62	63
Change in Production	0	0		Enter Decision	
Production	600	600	600	600	
Sales (dependent)	180	180			
Sales (independent)	420	423			
Sales (total)	0	603			
Change in Inventory	0	-3			
Inventory	0	-3			
Cost (Prod. Change)	0	0	= 2 * (0*0)		
Cost (Inventory)	0	9	= 1 * (-3*-3)		
Cost (Total)	0	9			
Accumulated Cost	0	9			

Conditions for current game: • 10 production units cause 3 sales units.
• Production is delayed by 2 weeks.

Figure 2.3. Computer screen display

Two emulated screen practice rounds were given in order to familiarize the subject with the system and their task. Each practice round allowed for one production decision. The practice rounds consisted of an actual screen shot on paper, (see Figure 2.3.), and two hypothetical examples. Each subject was informed that in the real experiment the computer would compute the consequences of the decision, but that for practice purposes the subject would perform the computations. Then, the subject was encouraged to make a decision for production and compute the consequences of the decision in long hand. When they had reached the line exogenous sales, the experimenter provided the subject with the number for exogenous sales. This number was the same for all subjects and determined in advance. At that time, the instructor took the opportunity to explain the nature of the random process that the computer used to generate exogenous sales. The subject then proceeded to compute the rest of the variables. A wrongly computed variable was immediately corrected and an explanation provided. Most of the subjects did not need any correction when computing the second practice task. This training process was utilized in order to provide subjects with the maximum amount of learning that could be attained, prior to turning on the system.

After the practice task, the instructor repeated the rules of the experiment. Subjects were informed that they would play 15 games altogether of an unspecified length. They were told that the conditions regarding delay of feedback, as indicated by delay of changes in production, and strength of feedback, as indicated in dependent sales, might vary from game to game. They were further informed that some conditions were intrinsically more difficult than others and were told to not be discouraged by what they might consider extremely bad scores. Specifics about what a bad score might be were not provided.

Subjects were informed that for the computation of their reward only the 12 best games would count, with the worst 3 discarded. However, all data were utilized in statistical analysis. Subjects were informed that the games would be played in four sessions altogether, three in the first session and four in the sessions to follow.

The computer was then turned on and the subject was asked if he or she had ever used a Macintosh computer before; if not, these mechanics were explained. The mechanics of the game were explained. In prepared note-pads, subjects were asked to write down their production decisions for each round of play on the sheet corresponding to the current round and ~~any calculations that they needed to make.~~ In this way, it was possible to monitor subjects' written decision processes unobtrusively and still provide the subject with a medium for calculations.

At the end of the session, the subject handed the note-pad back to the instructor and arranged an appointment for the next session. The subjects took about two to three weeks to complete all four sessions. To re-familiarize the subjects with the task, the instructor started each follow-up session with a quick summary of the objective of the task and the rules of the game including the random nature of exogenous sales.

2.7. Hypotheses

Both delay and gain modify and add complexity to the essential control loop as depicted in Figure 1.1 (and within the dashed line of Figures 2.4 and 2.5). Delay as implemented in the experiment modifies the action feed-out link by introducing a delay between 'Corrective action initiated' and 'Corrective action' (Figure 2.4). Gain as implemented in the experiment modifies the action feed-out link by linking unintended side effects to the initiated corrective action (Figure 2.5). How well decision-makers perceive the effects of the two complexifying factors and how well they adjust their corrective action to take them into account is the subject of this study.

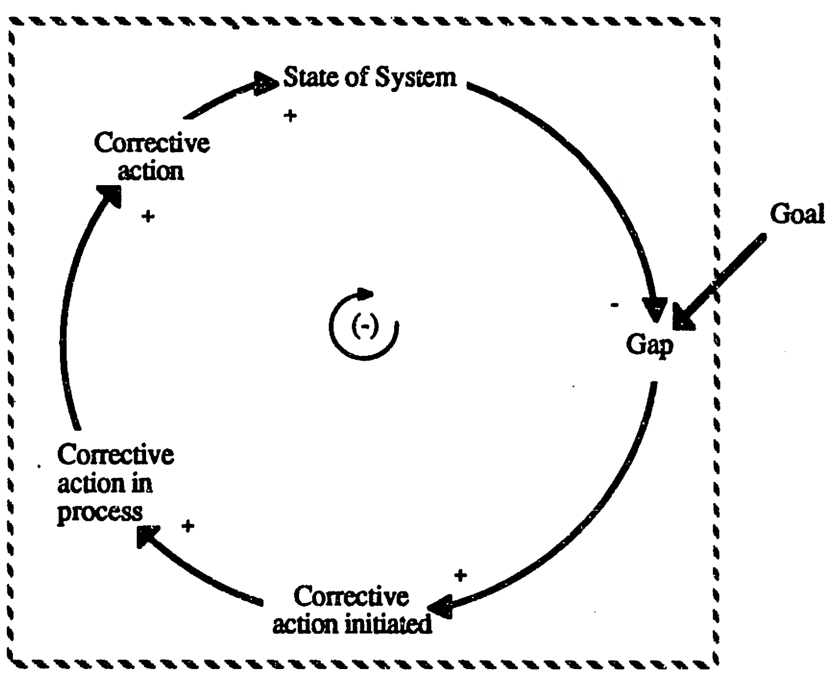

Figure 2.4. Closed-loop diagram of delay: Delays exist between initiation of corrective actions and their effects on the system

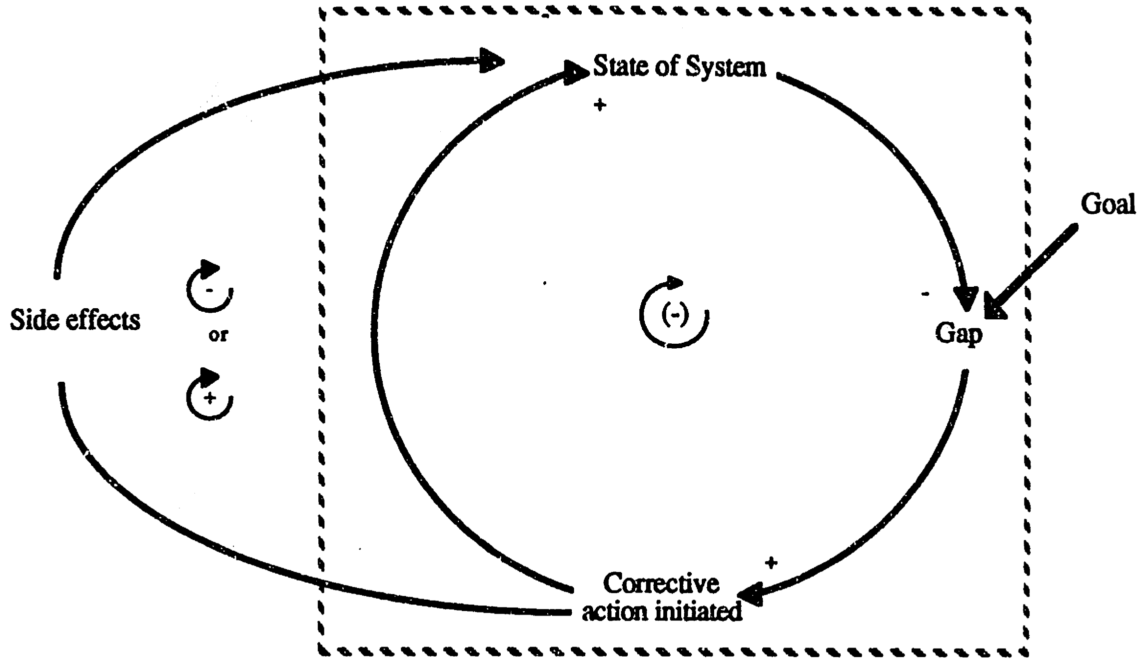

Figure 2.5. Closed-loop diagram of gain: intended and unintended side-effects vary in strength and direction

Decision models proposed in the literature can be characterized in terms of the assumptions they make about the number of links considered by the decision maker and the way that information is processed. Three rival models will be tested with that goal in mind: the rational model and two heuristic based models. Rational decision rules assume that the decision maker incorporates all of the relevant available information into his or her decision strategies. Heuristic decision rules assume that the decision maker utilizes a subset of the information upon which to base their strategy.

The rational model

Many researchers in the field of economics deem the optimal solutions to decision problems to be end-products of a long process of adaptation and market-induced selection in which the optimal solution is that decision rule that has evolved or survived over time and is now in a permanent state so long as the market or environment is at equilibrium. Lucas (1986,p. 402) explains:

> "I think of economics as studying decision rules that are steady states of some adaptive process, decision rules that are found to work over a range of situations and hence are no longer to be revised appreciably as more experience accumulates. From this point of view, the question whether people are in general 'rational' or 'adaptive' does not seem to be worth arguing over. Which of these answers is most useful will depend on the situations in which we are trying to predict behavior and on the experiences the people in question have had with such situations. It would be useful, though, if we could say something in a general way about the characteristics of social science prediction problems where models emphasizing adaptive aspects of behavior are likely to be successful versus those where the nonadaptive or equilibrium models of economic theory are more promising."

Hence, rational models may provide some useful information, but must be regarded only as a decision rule that people may aspire to after considerable experience, at best. Although normative models may indicate the end product of adaptation, they provide little understanding of decision rules used while systems are in disequilibrium or while agents are still learning and adapting.

The heuristic models

Behavioral decision theory, on the other hand, suggests that decision makers select the decision heuristics they use with respect to their *perception* of the requirements of the task to be solved. With regard to adaptation and heuristics, Tversky and Kahneman (1974) and Hogarth (1981) provide the impetus for the exploration of how people do in fact process the information in a dynamic task.

Previous studies on dynamic decision making suggest that subjects have the cognitive resources necessary to perform effective control; however, subjects fail to adapt to the presence of critical environmental factors, notably quality of feedback as measured in length of delay, certainty, etc. (Kleinmuntz, and Thomas, 1987; Payne, 1982; Simon, 1982; Tversky, and Kahneman, 1986).

Two heuristic models can be distinguished from one another in the amount of information they assume enters the decision process. The first model, the weight adjustment only model, assumes that people only examine information regarding the specific system variable to be controlled, for instance, the inventory gap, which they know should remain at zero. Weight adjusters therefore would ignore supply-line information and focus only on immediately decreasing the discrepancy between the desired goal state and the actual goal state of the system variable to be controlled. The weight and cue adjustment model, on the other hand, assumes that decision makers do incorporate supply-line information as well as goal discrepancy information. In contrast to the rational model, however, the weight and cue adjustment model does not utilize the complete information set, nor need it imply the optimal cue weights.

In this thesis, I will test which of these models most closely describes decision making as actually done by subjects. In addition, since each model leads to different predictions with regard to (a) decision time spent at the beginning of each trial for the first decision and (b) decision time spent during each trial for later decisions and (c) score, I will use these data in order to distinguish between the three models.

Rational model- predictions [1]

The optimal rule utilizes all of the information cues available and minimizes expected costs in the stock-adjustment task at any given point in the decision sequence. As delay increases, additional information states are introduced which require additional calculation of weight coefficients. Feedback strength modifies the value of information by entering into the calculations of the weight

[1] Different possibilities exist in how people might make decisions consistent with the rational model. The predictions depend on the assumed path that decision makers take to arrive at optimum. For purposes of sharp contrast, predictions in this section are based on the extreme assumption that people derive the optimal decision by solving the appropriate mathematical equation system. The predictions have to be qualified accordingly if one takes a less extreme approach (Lucas, 1986).

coefficients, but does not add additional information states. The optimal solution has the form:

$$\text{Control action} = \sum W_i S_i \quad i = \text{state } 1,2,3,\ldots n$$

Each period of delay adds a state variable and a cue to the task. The optimal values of the parameters W_i are computed by solving the Riccatti equation (D'Azzo, and Houpis, 1981). The complexity of the Riccatti solution roughly increases with the square of the number of states. Thus, the computational complexity of the solution is merely a function of the delay length, not of the feedback strength. Therefore, the rational model would suggest that decision time should not vary with respect to feedback strength, but should increase quadratically with increases in delay length. Because the model incorporates all of the information, performance would be optimal under all conditions. The predictions for the optimal solution, compared to the zero delay and zero feedback strength condition, include:

1) Time spent on first decision:

a) Subjects will spend a substantial amount of time for the first decision, solving a complex, multi-order Riccatti equation. Time spent on the first decision will increase quadratically as delay length increases, since the matrix to be solved requires additional states.

b) Feedback strength, on the other hand, does not have an influence on time spent on the first decision, because it does not change the order of the matrix of the Riccatti solution; however, a feedback strength of zero might be the possible exception, since some of the calculations simplify.

2) Time spent on later decisions:

a) The optimal rule requires subjects to multiply each information cue by a weight computed at the beginning of the trial and add the results together. The computational effort required is obviously significantly less than solving a complex, high-order matrix equation. Thus, time spent on later decisions will be significantly less than for the first decision. The more delay states present require more multiplications to be executed by the subjects. Thus we would expect to see more time spent on longer delay conditions.

b) Feedback strength, on the other hand, will not change the computational effort needed to derive a decision, thus no effect of feedback strength is expected.

3) Scores:

All scores will be optimal.

Weight adjustment only model- predictions

This model suggests that subjects focus solely on the quantity to be controlled, i.e. inventory gaps. Inventory is the variable which generates performance, thus the most salient. Furthermore, it is the only one for which a reference value can easily be obtained. The inventory gap should be 0. In contrast, it is uncertain as to what the appropriate values for production or sales should be at any moment in time.

The 'Weight Adjustment Only' model would suggest that decision time does neither vary with delay time nor with feedback strength. Hence, performance is best in a no-feedback situation and gets worse with increased feedback strength and increased delay time. Specific predictions for this model include:

1) Time spent on the first decision:

a) Since calculations are presumed to be done intuitively, time spent on initial decisions is shorter than the rational model.

b)Feedback strength has no effect.

2) Time spent on later decisions:

a) Delay has no effect.

b) Feedback strength as well has no effect.

3) Scores:

a) Increasing delays cause significantly increased costs relative to optimal, since this model neglects all the information that additional delay states introduce.

b) Increasing feedback strength also increases costs relative to optimal, since it is assumed that weights are not sufficiently adjusted in response to increases in gain.

Weight and cue adjustment model- predictions

This model suggests that subjects take additional information into account besides observing the variable to be controlled. The "Weight and Cue Adjustment" model provides for the possibility that subjects consider additional information that is generated as a consequence of increasing the number of delay states. However, they do not take all of the information into account as the rational model assumes. The way the information is combined and processed is based on implicit heuristics rather than the notion of an optimizer strategy.

The 'Weight and Cue Adjustment' model would suggest that decision time varies linearly with delay and is invariant with respect to feedback strength. Relative performance is best under a no-feedback situation and gets worse with both increased positive feedback strength and delay, assuming sub-optimal weights. The weight and cue model compared to the weight only model makes the same predictions about increasing feedback strengths, but their differences are apparent with increasing delay. Specific predictions include:

1) Time spent on the first decision:

a) Longer discrete delays imply a longer first decision time because subjects must deal with additional information.

b) Feedback strength, conversely, does not influence time spent.

2) Time spent on later decisions:

a) Should increase with increasing delay as additional multiplications have to be performed.

b) Feedback strength has no effect

3) Scores:

a) With increasing delays, score worsens, but not as severely as the "weights only model," since this model incorporates increasing delays to a degree.

b) Increasing feedback strength causes score to worsen but not significantly compared to delay conditions and Weights only model.

Figures 2.7.a, b, & c summarize the predictions made by each model over all of the levels of each variable. While some models make the same predictions, there are distinct differences in several of the conditions. I will use those differences in the results section to distinguish between the models.

Rational Model

	Increasing Delay	Increasing Gain
Score relative to optimal	score always equal to optimal	score always equal to optimal
Time Spent on First Decisions	increases approximately as the square of the number of states	no effect, with the possible exception of zero gain
Time Spent on Later Decisions	increases	remains the same

Figure 2.6a Rational model: treatment predictions

Weight Adjustment Only Model

	Increasing Delay	Increasing Gain
Score relative to optimal	worsens	worsens
Time Spent on First Decisions	remains the same	remains the same
Time Spent on Later Decisions	remains the same	remains the same

Figure 2.Cb Weight adjustment only model: treatment predictions

Weight and Cue Adjustment Model

	Increasing Delay	Increasing Gain
Score relative to optimal	worsens	worsens
Time Spent on First Decisions	increases	remains the same
Time Spent on Later Decisions	increases	remains the same

Figure 2.6c Weight and cue adjustment model: treatment predictions

3. Results

3.1 Overview

Throughout the analysis section, results will be repeatedly reported within the framework of a standard presentation "toolkit." Therefore, it is useful to fully understand the elements of that toolkit. We provide an introduction in the following. At the core of this section is the analysis of variance. All of the results are presented in the following order: 1) a mean table to provide the reader with a brief orientation as to what happened in the various conditions, 2) results of the ANOVA are presented along with the significance level of the treatment, practice, and subject effects, 3) each of these effects will be further illuminated with a figure showing the effects.

In the three results figures a "box and whiskers" template shows the median, longest, shortest, third longest, and third shortest dependent measure, as in Figure 3.1. The outlying horizontal tick marks represent the extreme scores for a subject, treatment, or sequence, while the rectangle defines the third largest and third smallest scores, and the median is shown as the horizontal tick mark located between the upper and lower ends of the rectangle.

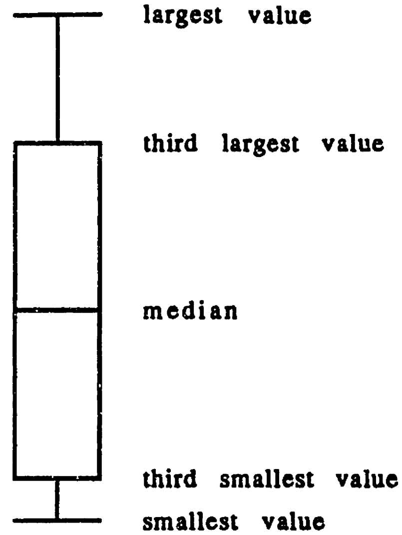

Figure 3.1 Box and whiskers template

Using the toolkit, we will investigate four different series of data: 1) analytical benchmarks to provide an understanding of the task itself and the ranges of possible performance, 2) scores actually achieved by subjects and their comparison against the various benchmarks, 3) time spent by subjects which is one indication of the heuristics used by the subjects, and at that time, we will start to complement the ANOVA's of the subject's effect with a regression analysis in which we attempt to investigate to what extent differences between "high" and "low" performance can be attributed to differences in the amount of time spent to make a decision, and 4) an analysis of control effort as measured by the amount of control exerted by the subjects over the system. Control effort analysis concludes the analysis of directly observable data supplied in chapter 3, then, in chapters 4 and 5, we will provide notebook and regression measures that allow further insight into the heuristics employed by subjects.

Each of the four sections will be presented within the same outline template: 1) introducing the measures utilized and the motivation for using them, 2) analyzing the data with the toolkit previously described, and 3) an overall summary for each subsection.

3.2 Benchmarks

In the following section, we will discuss the performance of the optimal rule, no-control rule, and the differences in performance between the two rules in order to provide an illumination of the controllability of the task as further described in subsection 3.2.3.

In describing the performance of subjects in the stock-adjustment task, it makes sense to provide an understanding of the possible range of performance. We need benchmarks to assess subjects' behavior. The benchmarks that I have chosen differ with respect to the level of information and processing that they utilize: the optimal control rule, defining the "ceiling" of performance, and the zero-control rule, providing a measure of the "floor." There are an infinite number of rules that fall between these two benchmarks, including subjects' rules.

A simulation analysis was conducted in parallel with the actual analysis. Fifteen "computer subjects" played through all of the trials in the same order and with the same exogenous sales as the fifteen actual subjects. Thus, there are 225 runs for the "optimal rule subjects" and 225 runs for the "no-control rule subjects."

3.2.1 Analysis of "ceiling"

As described in the experimental design section, the optimal control rule utilizes all of the information cues available and combines them in the best possible manner, i.e. the Ricatti solution. The optimal rule discards any information that is redundant and unnecessary. For example, the optimal rule ignores all "history" in the sense that past values of any state carry no information for the current value of the state. As far as the optimal rule is concerned, it does not matter what path the variables, such as inventory, took to reach their current state. The optimal rule utilizes information about the current state of the system in the form of current inventory, production, sales, and the various stages of supply-line production.

The optimal rule strives to bring the system towards a state of equilibrium. In equilibrium, inventory is zero and production equals sales. In equilibrium, no further production changes are indicated resulting in expected costs per round of zero, thus achieving minimum expected costs. It is important to note that the optimal rule does not have any more information than subjects do, and it must respond to the same exogenous random walk throughout a game. Following the optimal expectation for demand, the rule acts on the prediction that independent sales in the future will be the same as they were in the previous round.

Figure 3.2 illustrates the weights that the optimal rule associates with each variable under the treatment conditions of delay and gain. The optimum solution is not contingent on the particular random walk. Rather, it is general for any input (any random walk or any stationary input). The weights were derived numerically by using Matrix X as installed on the Athena computer system at MIT. Inventory^{o} , sales^{o}, and production^{o} indicate the current value of these variables. Production-1 to 4 indicates that these variables are future production states. The weights associated with each variable indicate the relative importance that the optimal rule attaches to that piece of information.

Weights Assigned to Information Cues
by the Optimal Rule

Delay	Cue	Gain -0.6	Gain -0.3	Gain 0	Gain +0.3	Gain +0.6
0	Inventory°	-0.3233	-0.3508	-0.3840	-0.4258	-0.4834
	Sales°	0.4943	0.5799	0.7051	0.9105	1.3316
	Production°	-0.4943	-0.5799	-0.7051	-0.9105	-1.3316
	Production1	0	0	0	0	0
	Production2	0	0	0	0	0
	Production3	0	0	0	0	0
	Production4	0	0	0	0	0
2	Inventory°	-0.3551	-0.3703	-0.3840	-0.3956	-0.4051
	Sales°	0.8640	1.0963	1.4731	2.1788	3.9512
	Production°	0	0	0	0	0
	Production1	-0.3551	-0.3703	-0.3840	-0.3956	-0.4051
	Production2	-0.5090	-0.7260	-1.0891	-1.7832	-3.5461
	Production3	0	0	0	0	0
	Production4	0	0	0	0	0
4	Inventory°	-0.3759	-0.3823	-0.3840	-0.3823	-0.3782
	Sales°	1.1121	1.5420	2.2410	3.5341	6.7349
	Production°	0	0	0	0	0
	Production1	-0.3759	-0.3823	-0.3840	-0.3823	-0.3782
	Production2	-0.2911	-0.3384	-0.3840	-0.4262	-0.4640
	Production3	-0.2166	-0.2990	-0.3840	-0.4663	-0.5428
	Production4	-0.2285	-0.5223	-1.0891	-2.2593	-5.3499

° Denotes current state of the variable

# Denotes future states of the variable

Figure 3.2 Optimal cue weights

Under zero delay, you will note that the weights for sales and production are identical. The optimal rule makes intuitive sense in that it strives for what would be an overall decrease in costs and, thereby, system equilibrium. In a zero cost equilibrium, two conditions need to be satisfied simultaneously. Inventory needs to be zero and production needs to equal sales. If only the first condition were fulfilled, production and sales would not balance out; inventory would begin to accumulate to the difference between the two rates leading to an inventory gap. Only when the weights on production and sales are exactly equal will the computed result for production change be zero, the only value that does not change production, and thus does not disturb equilibrium.

In terms of increasing delay, notice that the optimal rule does not place any significance upon current production when production delays exist, since current production does not hold information value in the presence of predetermined future production values. As seen in the weights, the importance of inventory remains relatively constant across gains, while the emphasis on sales and production increases quite dramatically across gain and delay conditions.

Negative gain aids in the control of the system while positive gain, a self-reinforcing loop, makes control more difficult. The optimal rule, concurring with our intuition, adjusts the weights given to sales and production in positive feedback conditions more dramatically. Likewise, comparing the ratio between inventory and the discrepancy between production and sales, the optimal rule puts more and more attention on the discrepancy as opposed to inventory. The relative importance of a rate discrepancy increases with increasing gain. Hence, the optimal rule shifts to a more "aggressive" derivative control.

Additionally, notice the weights associated with current inventory and production delayed by one round, P1; they are the same. The reason for this is that the optimal rule deems the one-round, delayed production as merely an indicator of the level of current inventory. For instance, if we were faced with ordering heating oil that would be shipped two weeks from now and if we had ordered heating oil previously that would arrive by truck shortly, then we count the oil in our tank and the oil on the truck in assessing our current inventory. Thus, it makes sense that the inventory in the tank and arriving by truck should be weighted equally when making further ordering decisions.

Overall performance

Costs diffe. widely among treatments. As gain and delay increase, so do optimal costs, indicating a dramatic difference in objective difficulty. In the most difficult task, gain +0.6, delay 4, costs are on average more than five hundred times as high as in the easiest conditions, gain -0.6, delay 0. In order to facilitate analysis and comparisons of subjects' costs and optimal costs, the $\log_2$ of the costs was used for further analysis. Figure 3.3 provides the mean $\log_2$(costs) per condition.

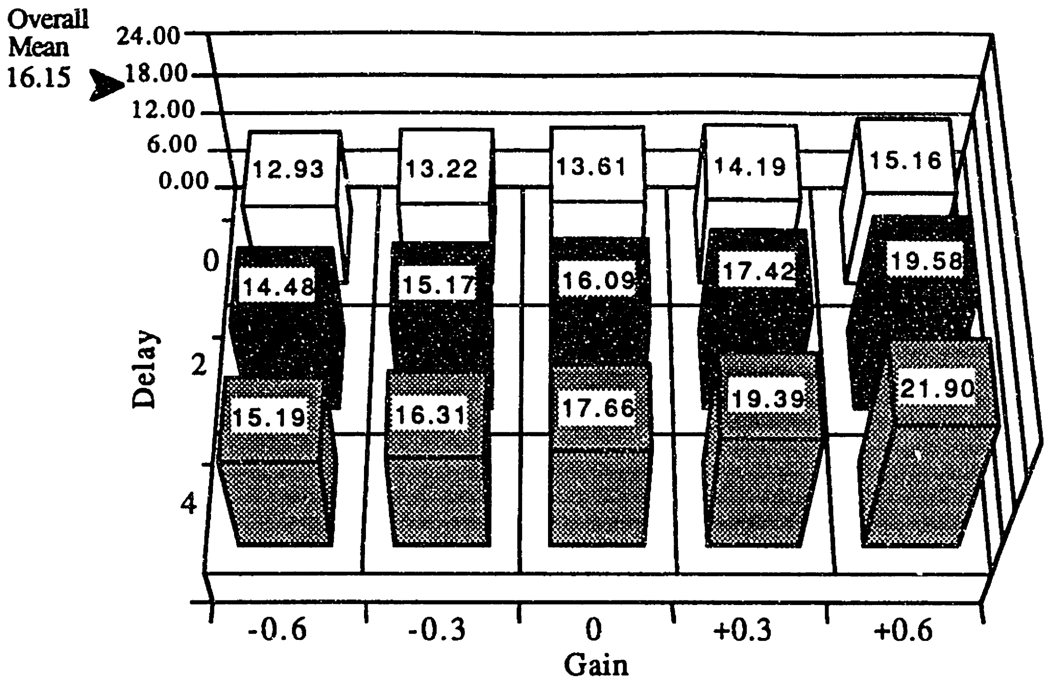

Figure 3.3 Optimal rule: mean costs in $\log_2(\$)$ across treatments

Analysis

Source	df	F	probability
Overall treatment	14,182	2432.55	p<.01
Delay	2,182	8449.03	p<.01
Gain	4,182	3708.37	p<.01
Delay x Gain	8,182	290.52	p<.01
Subject	14,182	1.08	-----
Practice	14,182	60.96	p<.01

-- not significant

Figure 3.4 Optimal rule: costs in $\log_2(\$)$ - statistics

In order to familiarize the reader with the standardized toolkit, used throughout the thesis, the computerized subjects' data will be analyzed with ANOVA in the same way the actual subjects' data are analyzed. We use the $\log_2$ of the costs instead of the actual costs as the dependent variable. As Figure 3.4 shows, not surprisingly there is no subject effect, since, after all, the "Macintosh subjects" are identical. Surprisingly, there is apparent evidence for a practice effect. However, the practice effect should be understood as a sequence effect arising from the 15 different random walks used across trials.[1]

Treatment effects

As we see in Figure 3.5, there are strong treatment effects with respect to delay and gain. Notice that the "box and whiskers" rise from the left corner to the upper right, indicating that the costs due to treatments increases steadily under increasing delay and gain. In addition, under increasing delay and gain there is an increase in variability.

[1] See the discussion in chapter 2 for a more detailed exploration of the relationship between the practice effect and random-walk patterns.

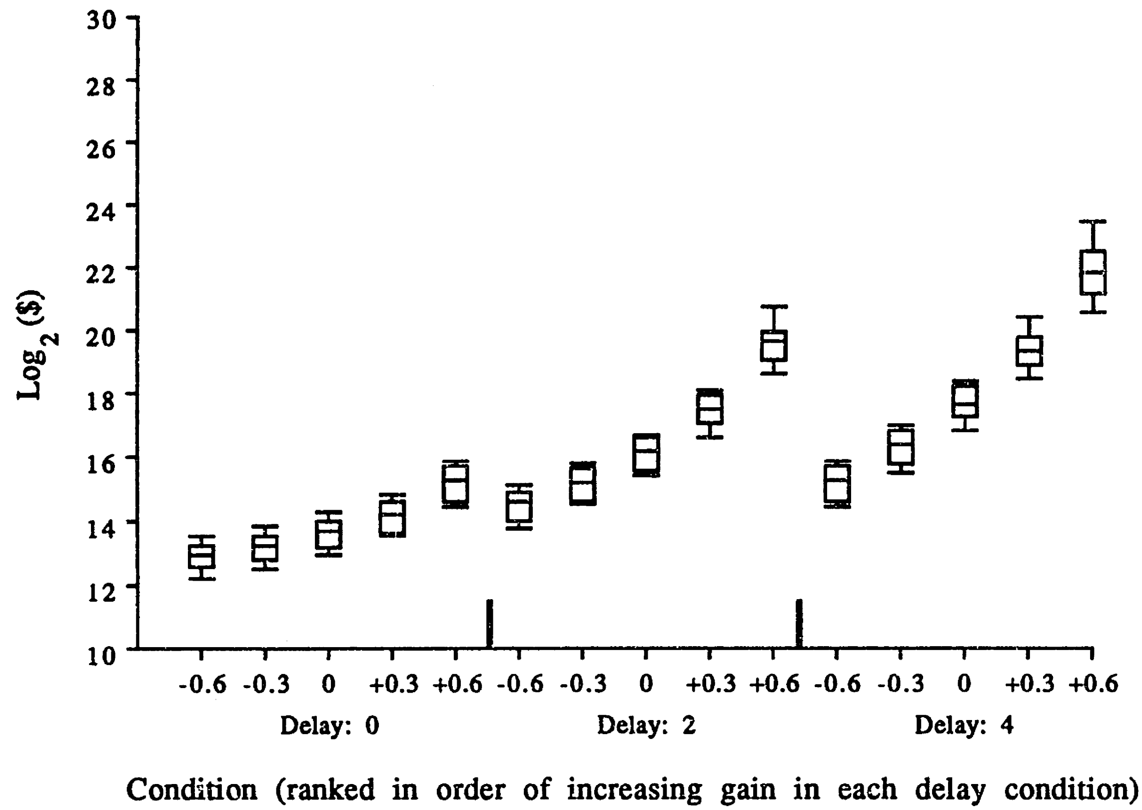

Figure 3.5 Optimal rule: costs in $\log_2(\$)$ - treatment effects

For a complete presentation of the toolkit, we provide the additional two figures, 3.6 and 3.7, showing practice and subject effects. The reader may peruse the figures in order to become familiar with the standard presentation toolkit. The subjects are listed with the numbers they were randomly assigned before the experiment, thus the label "pre-determined order." The trial numbers are presented with the session and game, e.g. 1.1 is session 1, game 1.

Practice effects

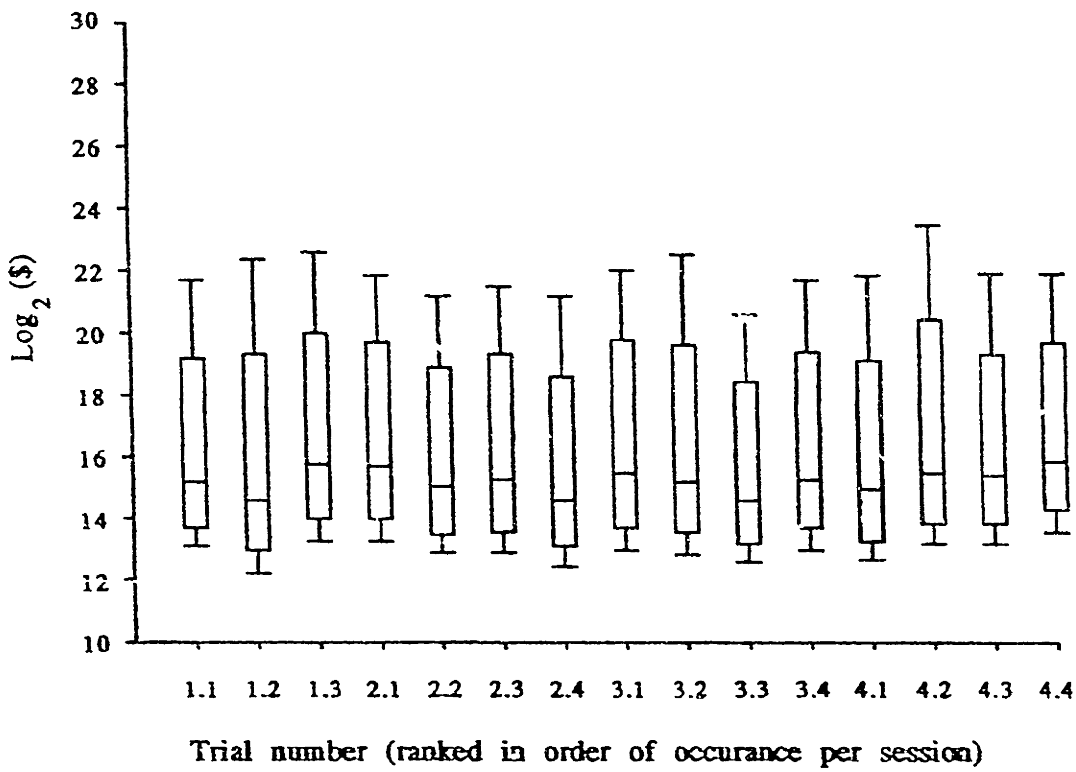

Figure 3.6 Optimal rule: costs in $\log_2$($) - sequence effects

Subject effects

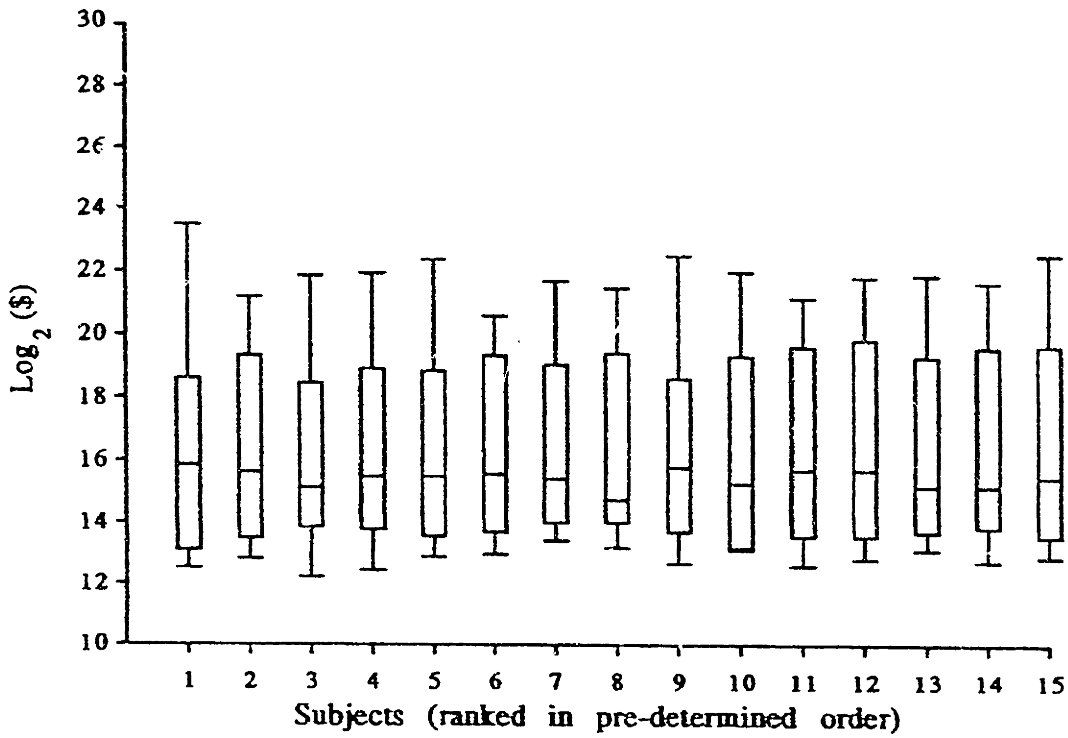

Figure 3.7 Optimal rule: costs in $\log_2$($) - "subject" effects - due to variations in random walks across sequences

3.2.2 Analysis of "Floor"

At the other extreme, the no-control rule, which serves as a "floor" for possible scores, operates by not performing any control whatsoever. Thus, changes in production are always zero and inventory is driven by exogenous sales alone. While the rule is extremely simple, it needs to be noted that the no-control rule is by no means the worst decision rule that subjects could follow. For example, overzealous control or poorly formulated control rule might very well destabilize the system rather than stabilize it and would do far more harm than simply refraining from any intervention at all.

Overall performance

Figure 3.8 shows the performance of the no-control rule. Average $\log_2$(costs) are the same for all conditions, remaining, on average, at 22.16 which is $9,174,292. Costs for all conditions are attributable to costs associated with inventory gaps that are driven by the random walk exogenous inputs. Also, with no changes in production, the side effects feedback loop is never active - gain makes no difference. Since the no-control rule does not make any production changes, it does not incur costs for changes in production. Thus, the costs are solely a function of the exogenous sales input to the system as they are accumulated into inventory. A random walk which does not wander far from its initial value will result in a smaller inventory discrepancy, and thus lower costs, while a random walk that wanders farther away from its initial value causes in each round huge accumulations in inventory. In fact, the most "difficult" random-walk input yielded costs 86 times as great as the "easiest" random-walk input.

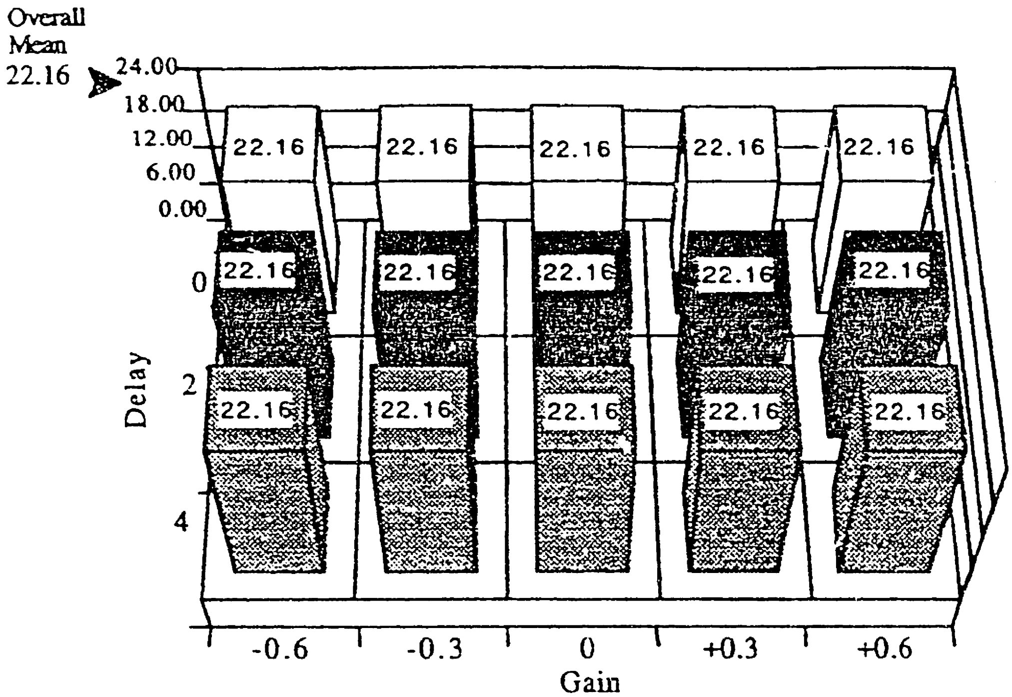

Figure 3.8 No-control rule: mean costs in $\log_2(\$)$ across treatments

Analysis

As shown in figure 3.9, there are neither treatment effects nor subject effects, as explained above. There are strong sequence effects. It is important to note, as we mentioned in chapter 2, that the practice effect is a combination of two things. Performance, as it changes over trials, is confounded with a particular instance of a random walk. These two effects cannot be distinguished from each other statistically.

Source	df	F	probability
Overall treatment	14,182	0.00	-----
Delay	2,182	0.00	-----
Gain	4,182	0.00	-----
Delay x Gain	8,182	0.00	-----
Subject	14,182	0.00	-----
Practice	14,182	infinite	p<.01

-- not significant

Figure 3.9 No-control rule: costs in $\log_2(\$)$ - statistics

The practice effect here, is dependent on the variability of a particular random walk. For completeness, Figures 3.10, 3.11, and 3.12 show what has just been described.

Treatment effects

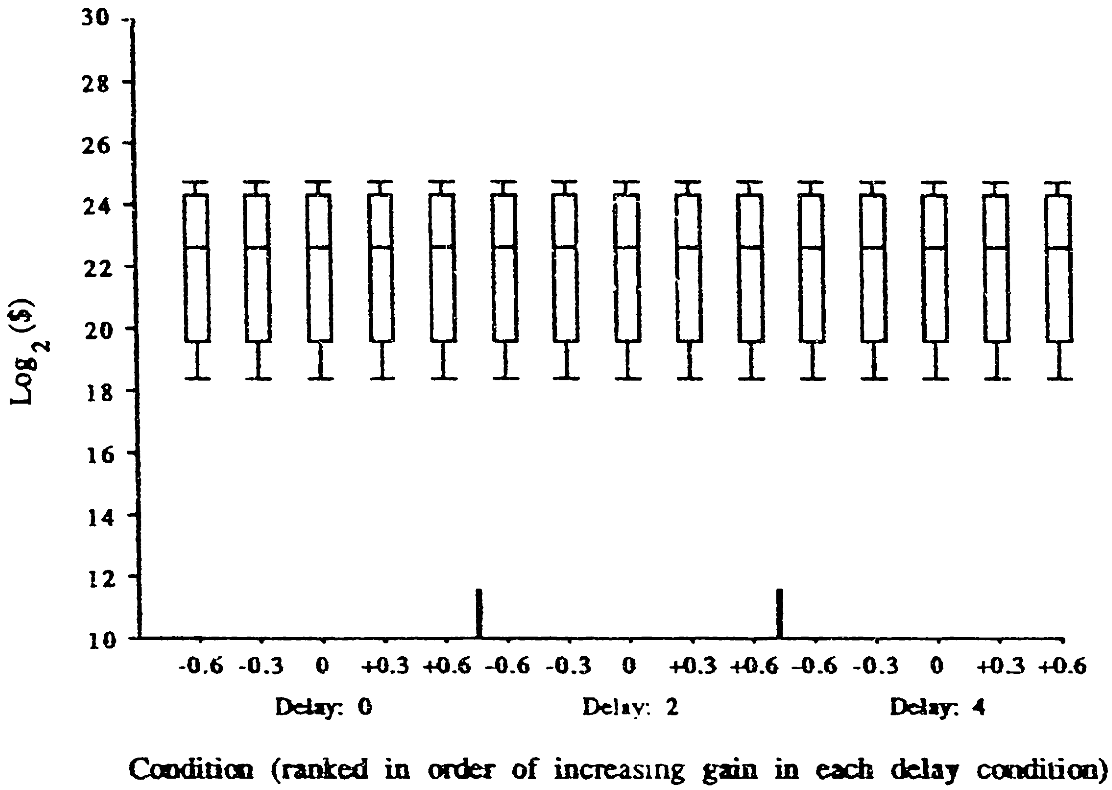

Figure 3.10 No-control rule: costs in $\log_2(\$)$ - treatment effects

Practice effects

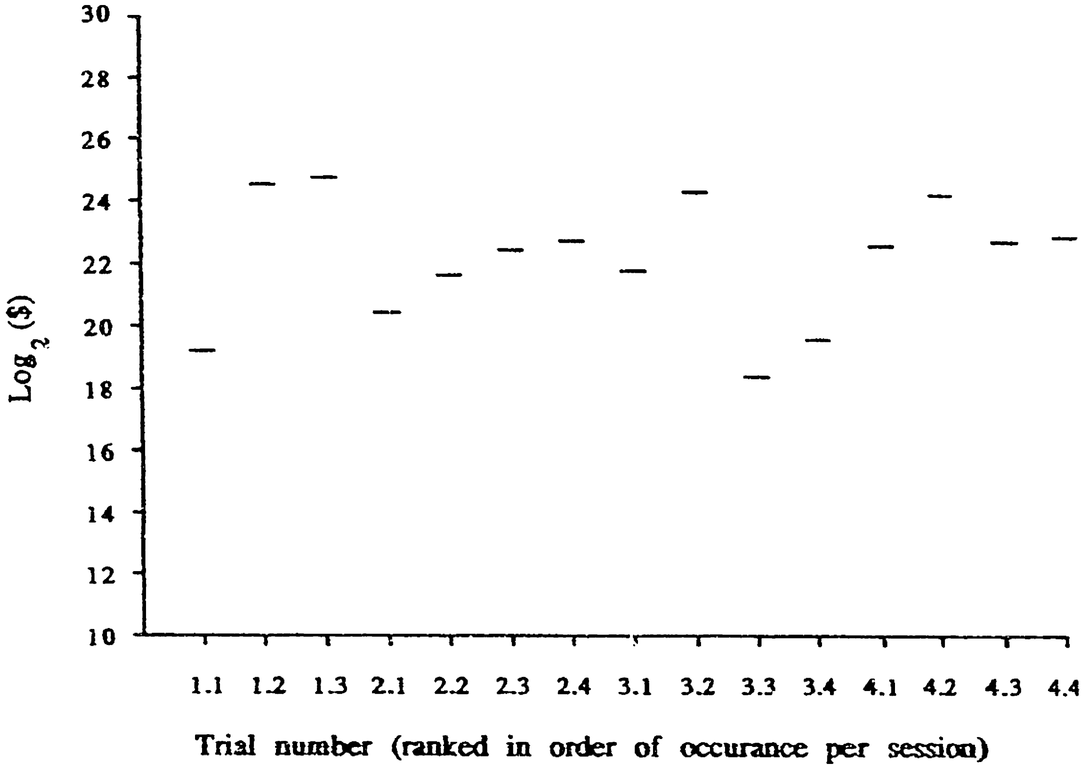

Figure 3.11 No-control rule: costs in log_2(\$) - practice effects

Subjec.: effects

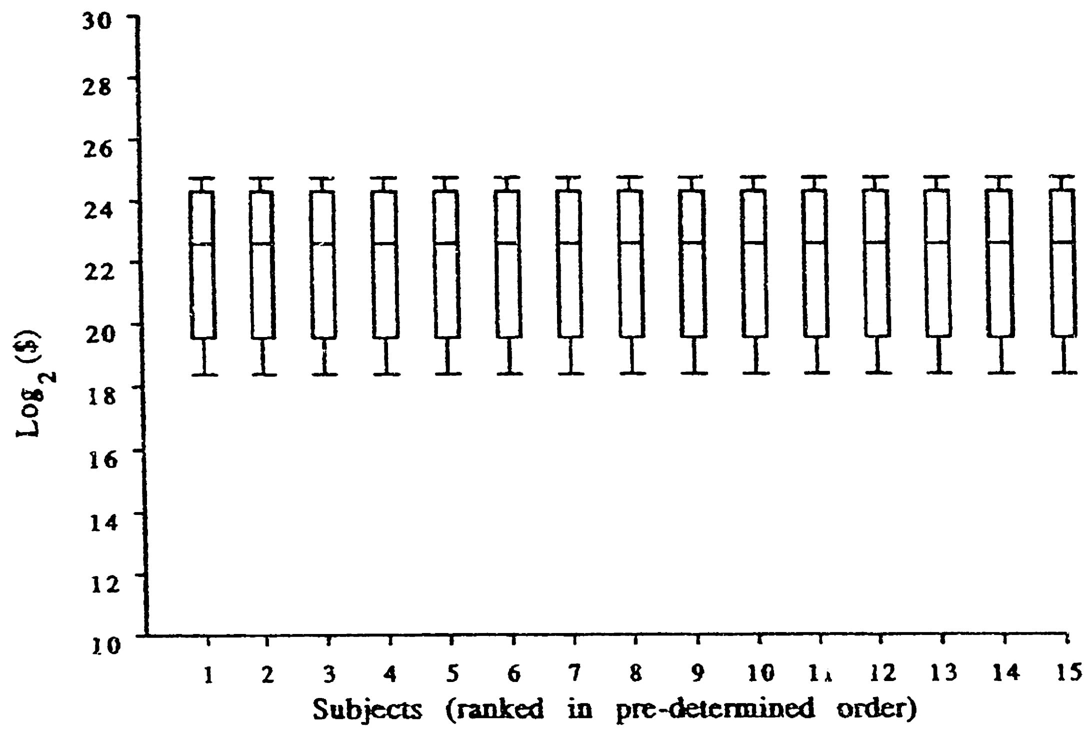

Figure 3.12 No-control rule: costs in $\log_2$($) - subjects effects

3.2.3. Analysis of "Ceiling" vs. "Floor"

Overall performance

The objective difficulty of the task is determined by the performance of the optimal rule, while the no-control rule determines a lower boundary of performance. Together, the two control rules define the region of controllability. As can be seen in Figure 3.13, the controllability region narrows considerably as difficulty increases. For instance, in the easiest condition, the optimal rule has costs 100 times lower than the no-control rule, and in the most difficult condition the no-control rule has only 20% higher costs than optimum. It is predicted that subjects' heuristics will produce performance somewhere within the region of controllability. However, the no-control rule may outperform a subject who uses a poorly formulated heuristic. Likewise, the optimum rule produces minimum expected costs conditional on the optimal expectation that future exogenous sales will equal their current value. However, a subject may get lucky in guessing which way sales will in fact move and may then outperform the optimal rule in any trial of finite duration.

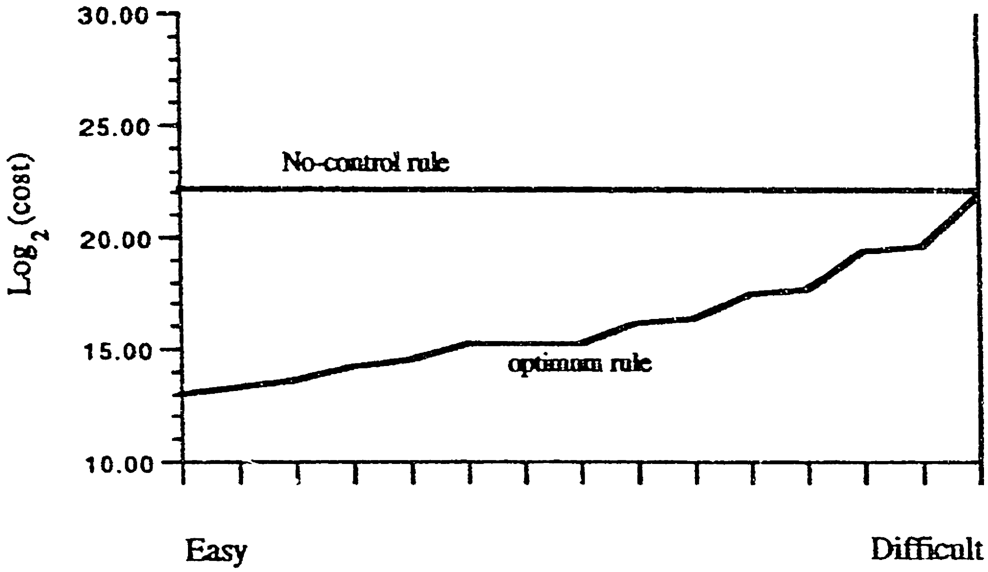

Figure 3.13 Controllability of the stock-adjusment task under increasing delay and gain

This phenomenon is illustrated even by the no-control rule. Compared to the optimum rule, the no-control rule gets "lucky" sometimes. Not only does the average of the no-control rule get close to the performance of the optimal rule, but the best of the no-control performances was better than optimum in the most difficult condition. As a matter of fact, the no-control rule achieves lower costs than optimum in five out of fifteen trials in the most difficult condition. The optimal rule expects exogenous sales to be exactly what they were in the previous round, and acts upon this belief with appropriate changes in production. The no-control rule, on the other hand, makes no production changes despite the varying random-walk patterns, and it therefore incurs only costs for inventory discrepancies and not for changes in production.

Whenever exogenous sales take a turn opposite in sign to the current inventory discrepancy, the no-control rule gets "lucky," as it sees the inventory discrepancy diminish. By contrast, the optimal rule is "caught by surprise" and is stuck with the consequences of unwarranted control actions. Obviously, by definition, the optimum rule is always better than the no-control rule on average.

3.2.4 Discussion of benchmark results

The result of dramatic differences between treatment conditions under optimal rule conditions underscores the need for explicit benchmarks. If we were to look at subjects' scores alone, we would confound objective difficulty of the task across treatment conditions with subjects' ability to deal with the different treatment effects.

The small gap between optimal and no-control performance for hard conditions seems to mean that the relative performance measure alone will not tell us how well subjects did. We will have to consider subjects' cue-weights and other data to determine how they behaved and not just what they did.

3.3. Subjects' Scores

3.3.1 Subjects vs. Benchmarks

To reiterate, the optimal control rule combines the necessary information in the best possible manner, while the no-control rule makes no control actions on the system, thus incurring larger costs. In terms of information cues combined, subjects' decision rules should fall somewhere between optimal and no-control. One could hypothesize that subjects operate as if they were in a no-delay, no-gain situation. In the presence of delay and gain, subjects adjust the weights that they attach to cues slightly to account for the presence of these factors. Subjects' scores, therefore, should on average fall somewhere between the upper and lower benchmarks.

Plotting subjects' average costs against average optimum and no-control rule's costs reveals a surprising result, shown in Figure 3.14. The subjects' average performance lies roughly parallel to the optimum performance, remaining approximately 4 times higher than optimum across the 15 conditions. The subjects' average costs even rise above the no-control rule's costs in the most difficult condition. Contrary to what one might have expected, subjects' costs are not an average of the two benchmarks, rather, they remain parallel but consistently higher than optimal.

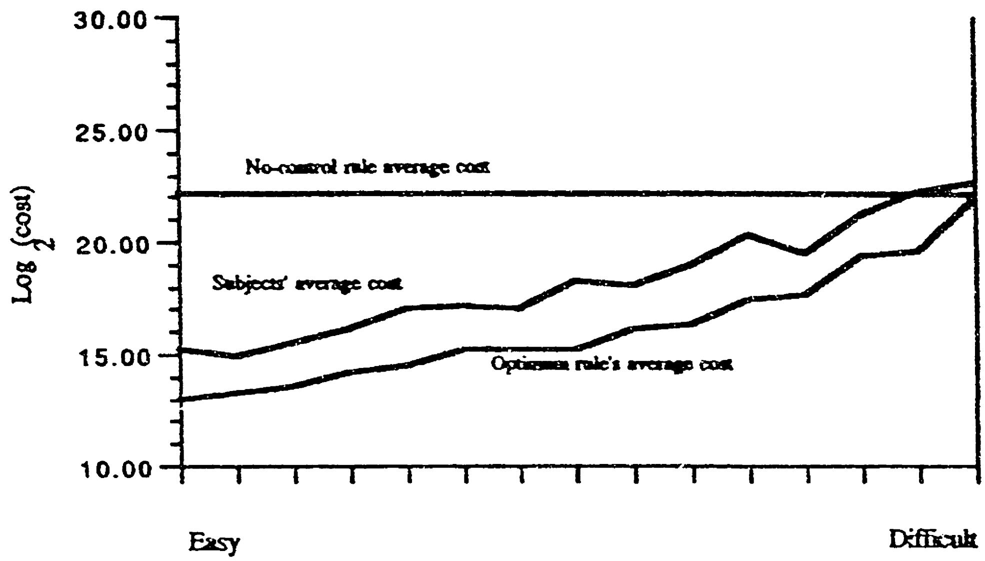

Figure 3.14 Controllability in $\log_2(\$)$: average optimum & average no-control vs. average subjects' costs

The question of whether to use the absolute difference between scores or the ratio remains open. I previously stressed the importance of benchmarks for comparing against subjects' performance and choose the optimum and no-control rules as a "ceiling" and "floor," but I have not yet addressed the issue of which manner of numerical comparison is most meaningful. Should we compare distances or ratios? Arguments can be made for both. Consider the following example. One subject achieves costs of \$20,000 in a condition in which the optimum rule results in \$10,000 and another subject receives costs of \$1,010,000 in a condition in which the optimum rule yields \$1,000,000. Did the subjects perform equally well, or did one perform better? Both subjects have \$10,000 higher costs than optimal, and after all, a dollar is a dollar. The first subject had costs twice as high as optimum, while subject 2 had costs of only 1% higher than optimum. It seems to us that the percentage differences give a stronger indication of performance. Thus, the majority of the analyses that follow were made on a percentage basis, but we provide the distance comparisons in section 3.3.3.

For the detailed presentation of subjects' performance, we will structure the discussion around five areas: 1) mean costs, 2) analysis of subjects' $\log_2$(costs), 3) absolute differences between subjects' scores and optimal scores, 4) ratios of subjects versus optimal, 5) ratios of subjects versus no-control.

The final costs achieved for each trial are used to evaluate overall performance for each subject. In accord with the rules set out in the methodology section, the three worst trials were discarded for each subject in order to determine just compensation for the experiment, but the results of all fifteen trials were used for the purpose of analysis as described in the rest of this section.

3.3.2 Analysis of subjects' $\log_2$(costs)

Overall performance

As shown in Figure 3.15, subjects' mean $\log_2$(cost) was 18.28. Costs range from 12.46 to 29.74. Again, we see differences due to treatments, suggesting that subjects were influenced by the presence of delay and gain. As expected, we see the worst performance under the most difficult conditions.

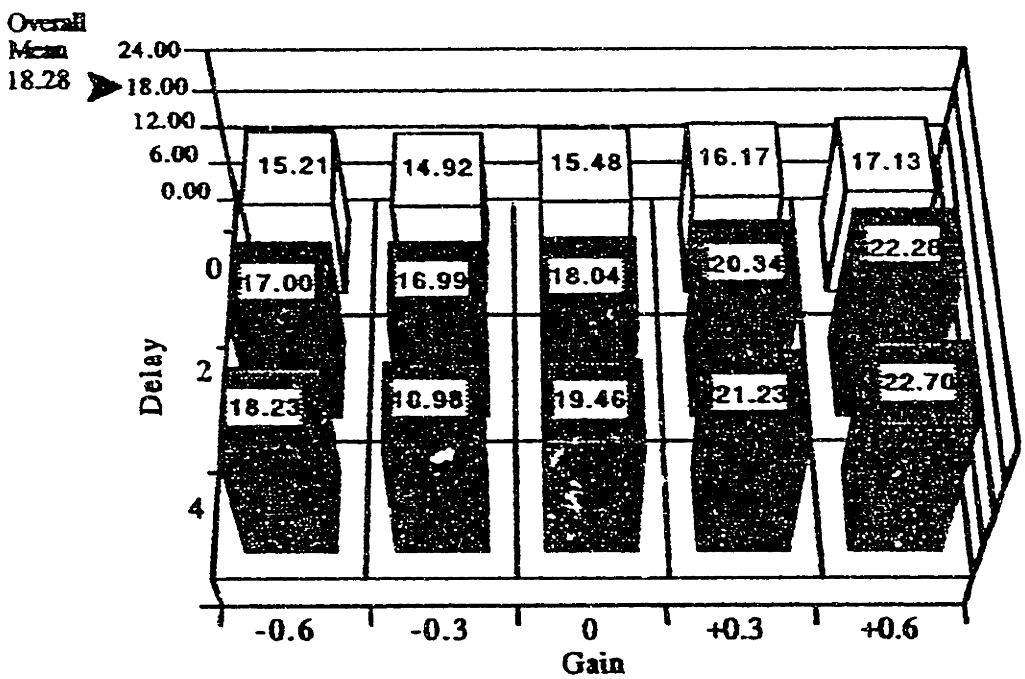

Figure 3.15 Subjects: costs in $\log_2$(\$) across treatments

Analysis

Source	df	F	probability
Overall treatment	14,182	24.65	p<.01
Delay	2,182	98.52	p<.01
Gain	4,182	32.53	p<.01
Delay x Gain	8,182	2.24	p<.05
Subject	14,182	5.57	p<.01
Practice	14,182	2.50	p<.01

Figure 3.16 Subjects: costs in $\log_2(\$)$- statistics

As Figure 3.16 reveals, there are significant effects of delay and gain (p<.01) and their interaction (p<.05). Likewise, subject and practice effects are significant (p<.01).

Treatment effects

Figure 3.17 shows subjects' $\log_2$(costs) for each of the 15 conditions. A strong effect of delay length and of gain is readily discernable. As both delay and gain increases, performance deteriorates.

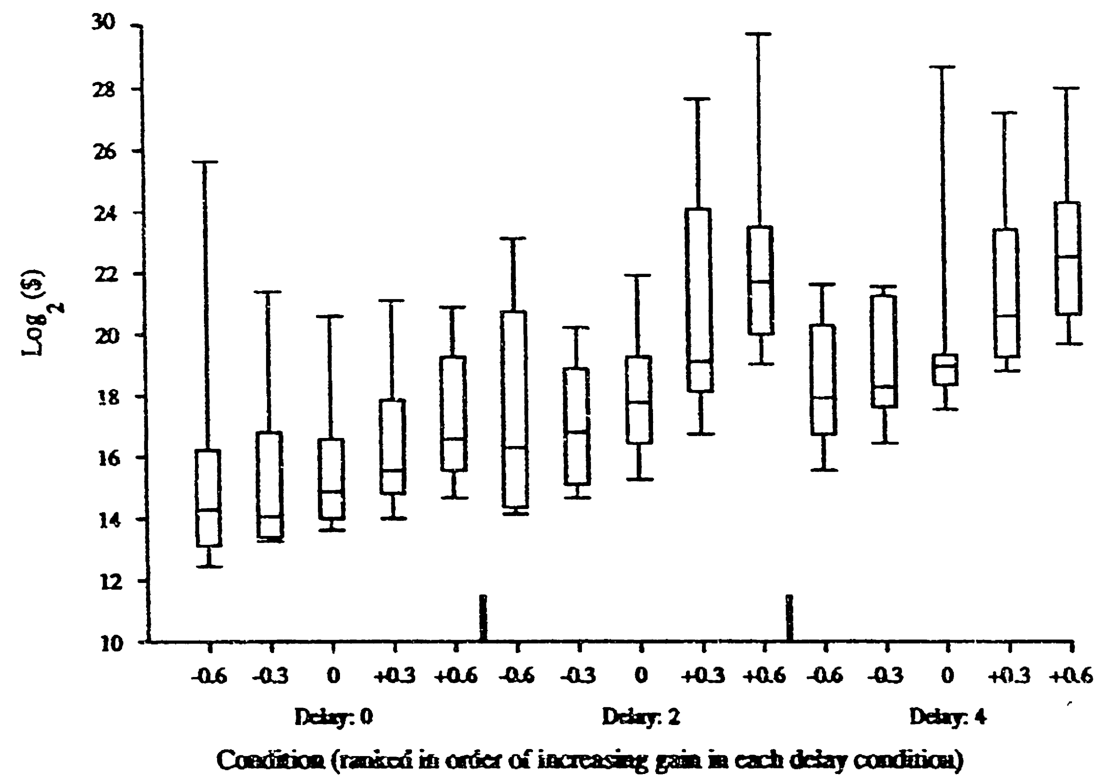

Figure 3.17 Subjects: costs in $\log_2$(\$) - treatment effects

Practice effects

ANOVA Figure 3.16 shows that there were differences between trials (p<.01). Figure 3.18 shows the subjects' score for each of the 15 trials of the experiment. Although the ANOVA revealed a significant effect of sequence, a clear pattern is difficult to see from the figure; however, it appears that performance in the first two trials contains greater variance and more poor performers.

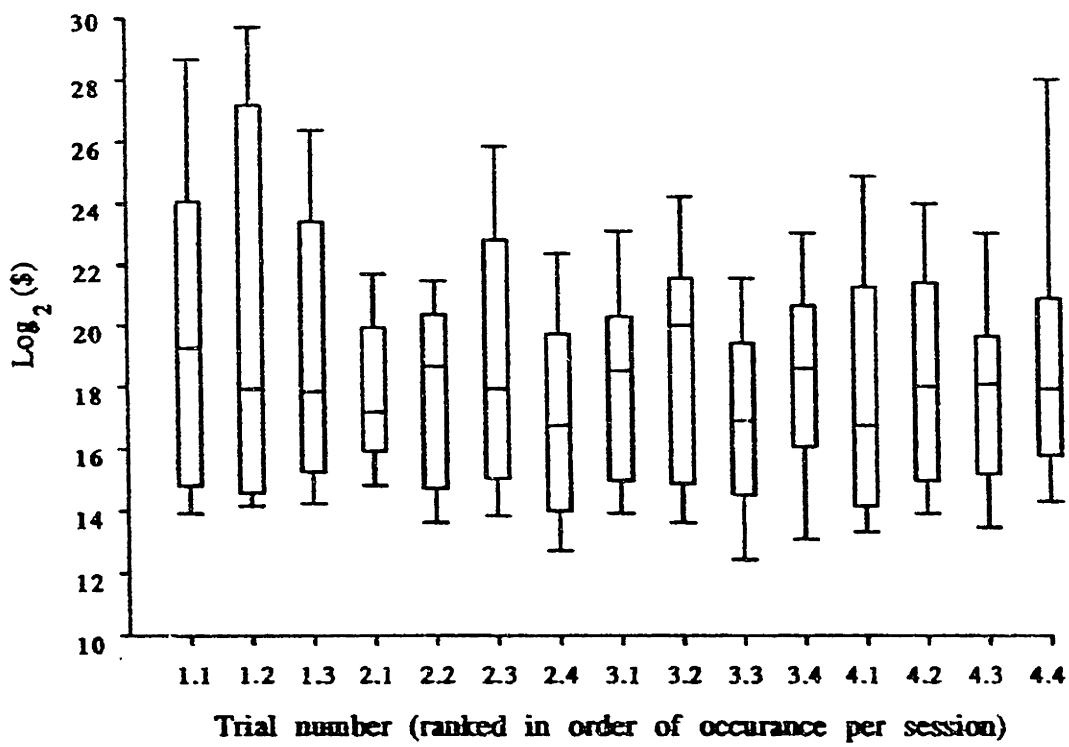

Figure 3.18 Subjects: costs in $\log_2(\$)$ - practice effects

Subject effects

There were significant differences between different subjects ($p<.01$), as Figure 3.16 reveals. Figure 3.19 gives an overview of subjects' performance. While the grand mean is 18.28, two subjects, (1 and 13), reach the grand mean only in their best trials, and their mean for all trials is approximately 20.

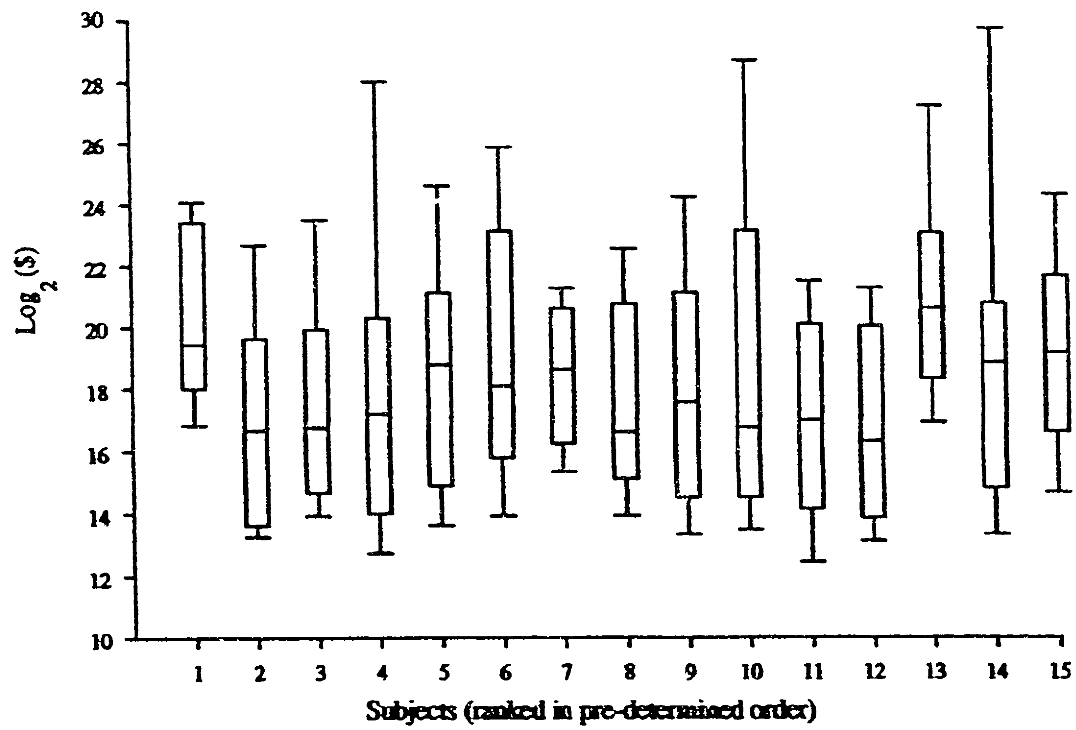

Figure 3.19 Subjects: costs in $\log_2$(\$) - subject effects

3.3.3 Absolute difference between subjects' scores and optimal

In order to show the difference between comparing absolute differences of performance and $\log_2$ differences of performance, the subjects' absolute differences in costs from optimum are supplied in Figure 3.20. The cost differences range from -2,858,306 to 893,364,683, while the overall mean is $10.7 million. Subjects show large increases in costs in the most difficult conditions, peaking in the 2 delay, +0.6 gain condition with a mean score of $68.24 million. The scores show an increased distance from optimum across conditions, except for the 0 delay, -0.6 gain condition. [2] Thus, we can see that there appears to be an effect of treatments.

[2] The mean in this condition is caused by an outlier of 51,979,418 which is considerably higher than the second highest score in that condition of 492,201.

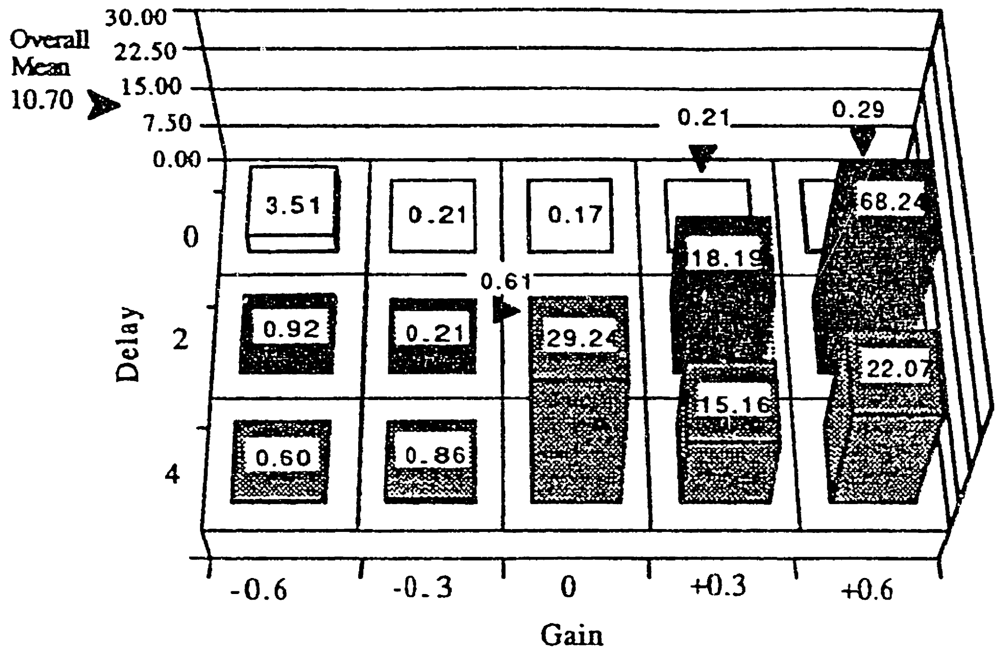

* scores presented in millions

Figure 3.20 Subjects vs. optimum: absolute differences in $

So far we have mainly looked at subjects' absolute costs as distance from optimum. What we've established was that subjects' performance deteriorates as the objective difficulty of the task increases. The following description provides a more detailed analysis of the $\log_2$ differences of subjects' performance compared to optimum and no-control rules.

3.3.4 Subjects vs. optimal

Overall performance

Figure 3.21 shows the average $\log_2$ difference scores for the comparison of subjects to optimal. On average, subjects have costs a little greater than four times as high as optimal, as indicated by the average difference of 2.12.

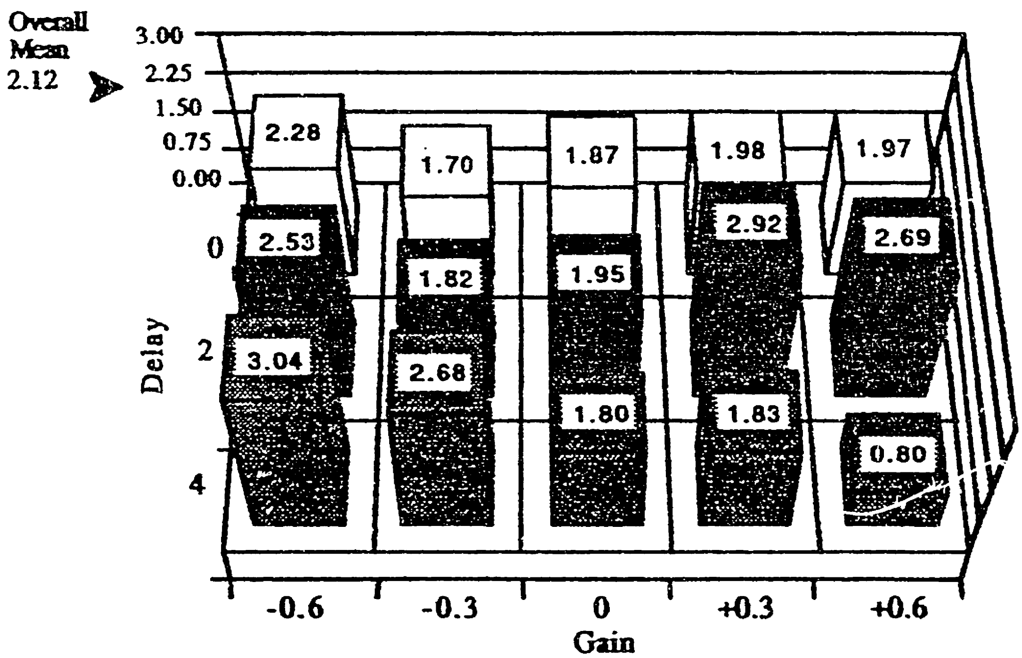

Figure 3.21 Subject vs. optimum costs: average $\log_2$ difference across treatments

It should be reiterated that applying the optimum rule is not an absolute guarantee for achieving the lowest costs possible for the reason that in computing the rule, the only knowledge used about exogenous sales is that they can be described as following a random walk. Being lucky and correctly anticipating the next values for exogenous sales would thus be one possible way to outperform the optimal rule. Even the no-control rule, for instance, achieves lower costs than optimum in five out of fifteen rounds in one condition, as mentioned in subsection 3.2.3.

Analysis

Figure 3.22 shows no significant influence of treatments, while subject and practice effects are significant ($p<.01$).

Source	df	F	probability
Overall treatment	14,182	1.31	-----
Delay	2,182	1.02	-----
Gain	4,182	1.22	-----
Delay x Gain	8,182	1.43	-----
Subject	14,182	5.55	$p<.01$
Practice	14,182	2.36	$p<.01$

-- not significant

Figure 3.22 Subject vs. optimum costs: $\log_2$ difference - statistics

Treatment effects

As the figure 3.22 revealed, there is no effect of treatments using the $\log_2$ differences.

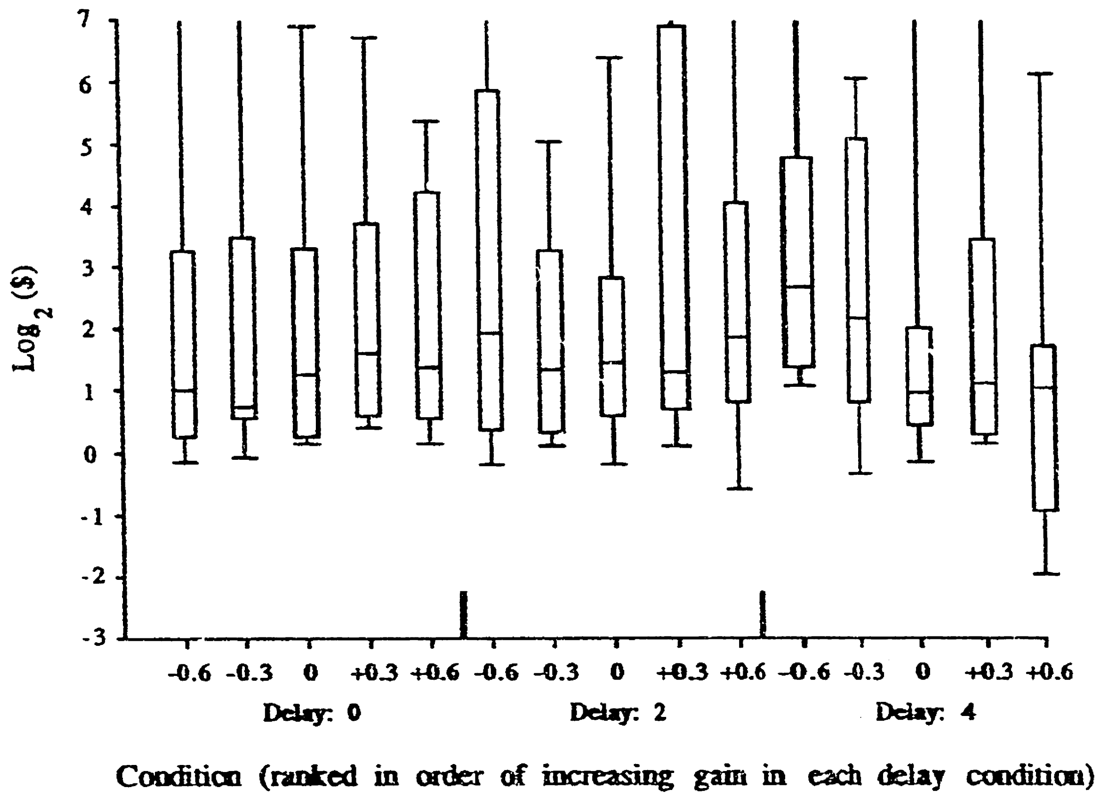

Figure 3.23 Subject vs. optimum costs: log2 difference - treatment effects

Practice effects

ANOVA Figure 3.22 shows a practice effect ($p<.01$), but Figure 3.24 shows that a clear pattern is difficult to discern. It seems, however, that the score is highest during the first two trials and continues to improve slightly during the rest of the trials. This trend is punctuated by several exceptions, however (Trial 2.3, 4.1).

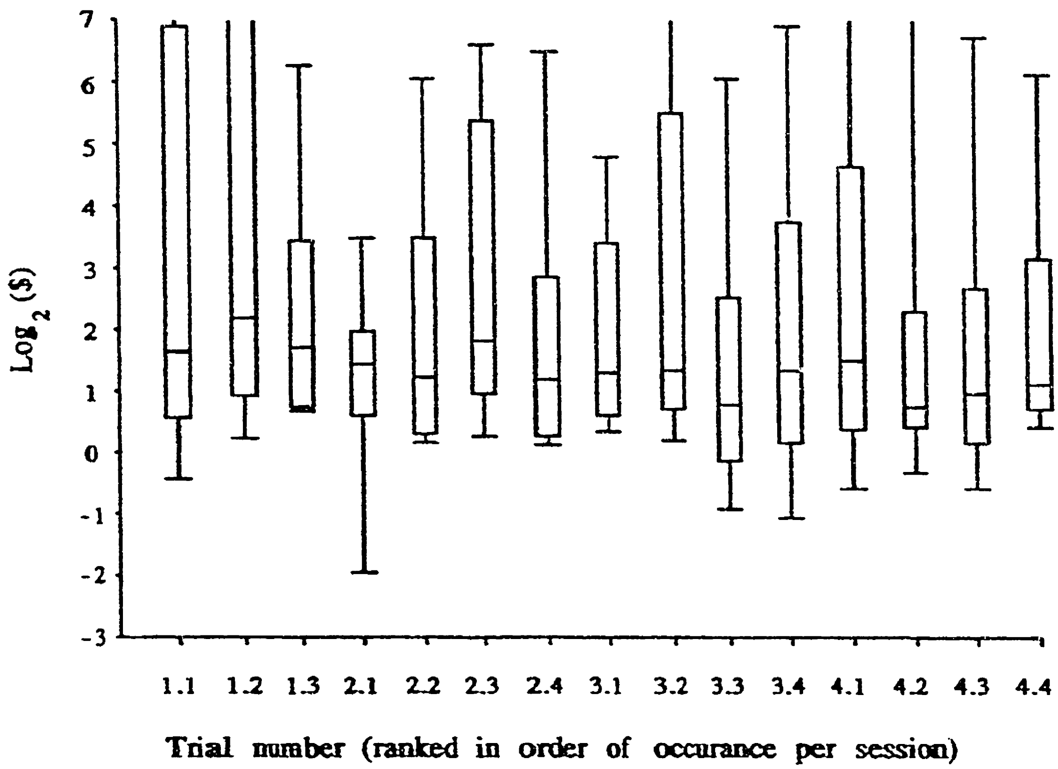

Figure 3.24 Subject vs. optimum costs: log2 difference - practice effects

Subject effects

Figure 3.22 reveals distinct difference among subjects ($p<.01$). Figure 3.25 reveals that, on average, the subjects end up with costs four times as high as those achieved by applying the optimal control rule. Subjects' performance varies considerably. While the lowest-performing subjects' median trial shows costs more than sixteen times higher than optimal (S1,S13), the highest-performing subjects produce costs less than twice as high as optimum in their median trial(S4,S11,S12).

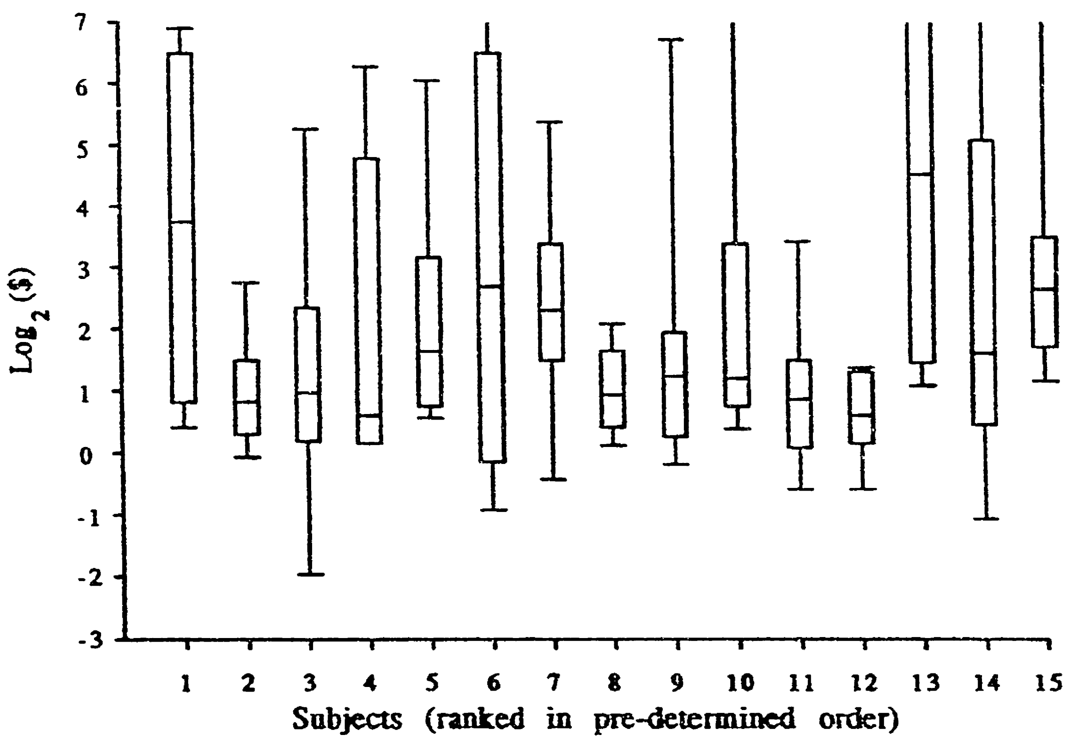

Figure 3.25 Subject vs. optimum costs: log2 difference - subject effects

3.3.5 Subjects vs. no-control

To compare subjects' performance against the benchmark set by the no-control rule, we employ the same $\log_2$ differences used in the previous sections.

Overall performance

Figure 3.26 shows the overall performance of subjects' costs versus no-control rule costs. Negative differences imply that the subjects achieved lower costs than the no-control rule. Overall, subjects achieve costs almost sixteen times lower than no-control. The difference of performance ranges from -10.705 to a maximum of 9.485, which implies that in some instances subjects perform more than 500 times worse than the no-control rule.

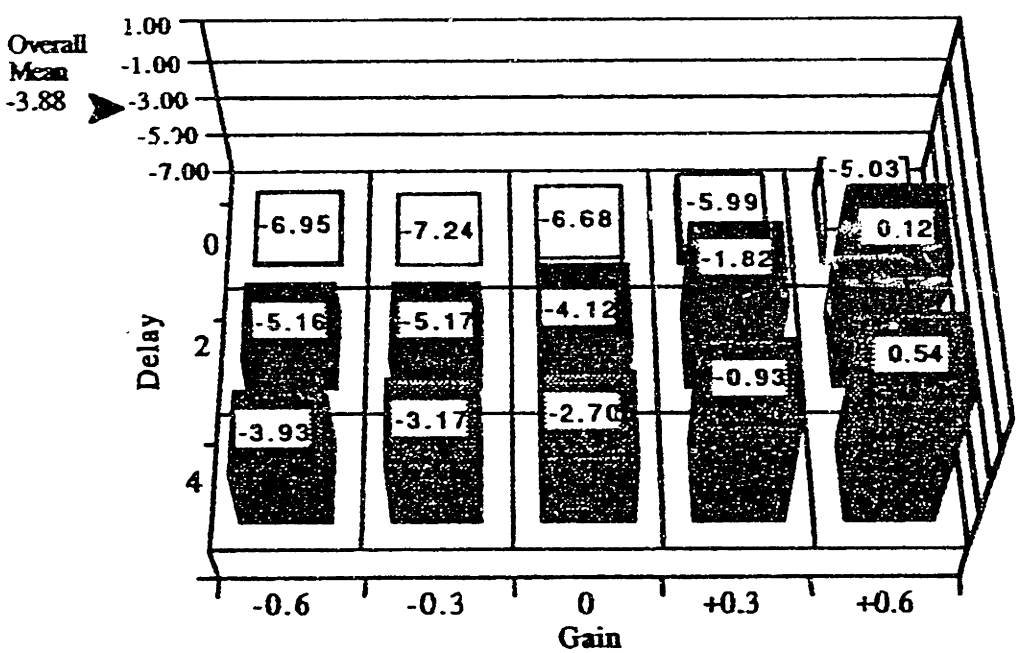

Figure 3.26 Subject vs. no-control costs: average $\log_2$ difference across treatments

Analysis

Source	df	F	probability
Overall treatment	14,182	24.65	p<.01
Delay	2,182	98.52	p<.01
Gain	4,182	32.53	p<.01
Delay x Gain	8,182	2.24	p<.05
Subject	14,182	5.57	p<.01
Practice	14,182	12.88	p<.01

-- not significant

Figure 3.27 Subject vs. no-control costs: $\log_2$ difference - statistics

Treatment effects

ANOVA Figure 3.27 shows that the overall effects of delay and gain were significant (p<.01), and the cross-treatment effect of delay length and gain is significant (p<.05). Figure 3.28 shows that performance differs widely for the different conditions.

While the majority of the subjects outperform the no-control rule on average by more than 123 times in the three most favorable conditions (Delay Length: 0; Feedback Strength: -0.6, -0.3, & 0), the no-control rule produces lower costs on average in the two least favorable conditions (Delay Length: 2 and 4; Gain: +0.6) by 1.3 times less than subjects. While the best subject outperforms the no-control rule in the most favorable condition (Delay Length: 0; Gain: -0.6) by more than 1000 times, the best subject in the least favorable condition (Delay Length: 4; Gain: +0.6) outperforms the no-control barely more than four times.

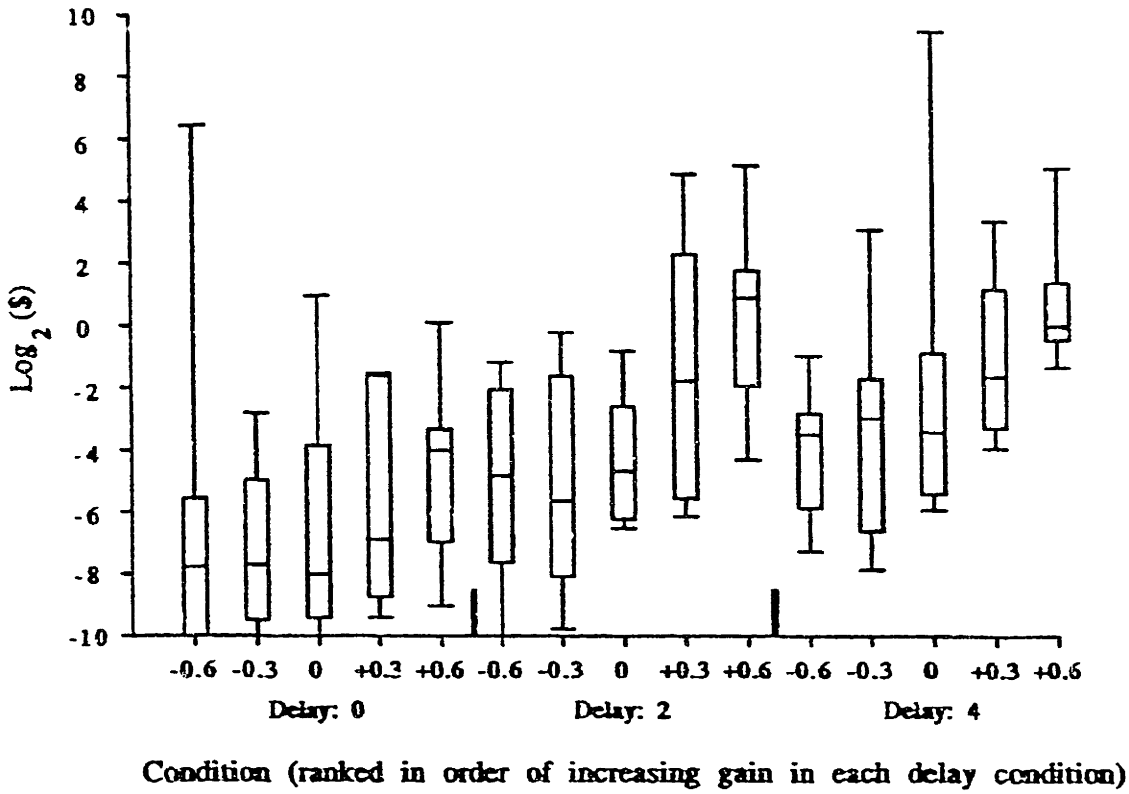

Figure 3.28 Subject vs. no-control costs: $\log_2$ difference - treatment effects

Two trends are readily apparent from the figure: 1. Relative to the no-control rule, performance deteriorates with increasing delay and 2. Performance decreases with increasing gain. In addition, it appears that performance is effected by the particular combination of delay length and gain: Long delay times and high positive gains cause a particularly steep decrease in performance.

Practice effects

ANOVA Figure 3.27 shows that the practice effect was significant (p<.01). Figure 3.29 displays that the no-control rule outperforms many of the subjects in their first trial. Subjects' performance in the rest of the trials is clearly improved. However, performance does not follow a recognizable trend, but is characterized by considerable variation. For example, subjects' performance falls unexpectedly in trial 3.3 and 3.4.

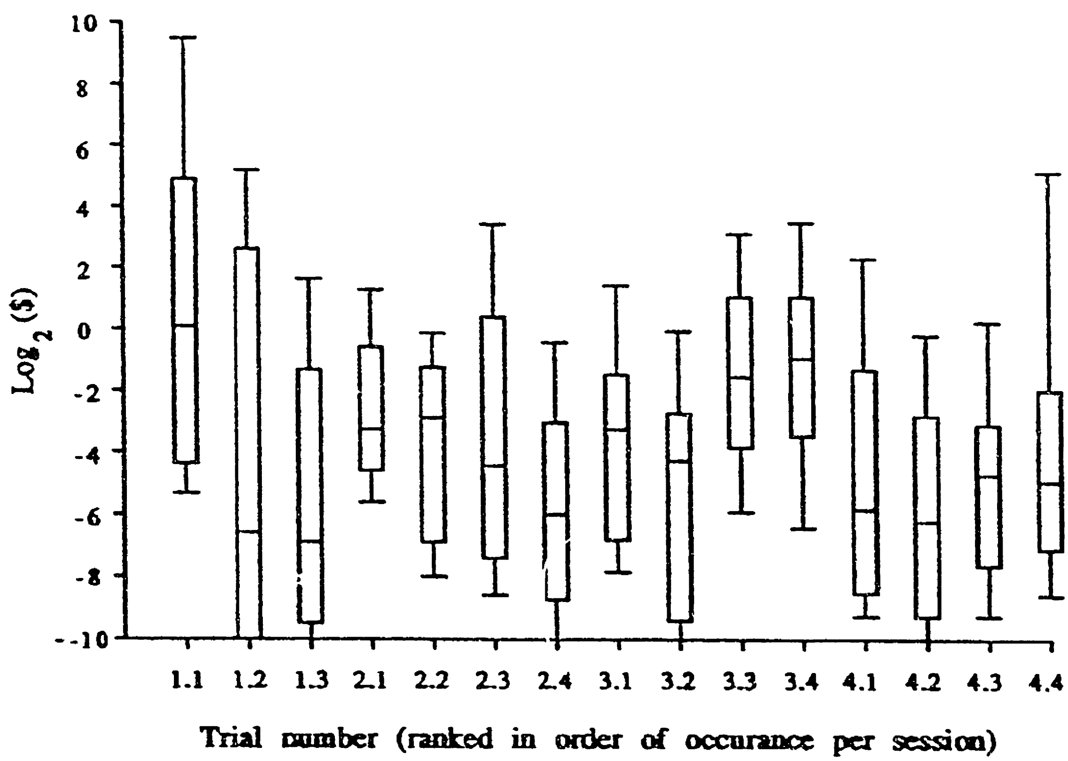

Figure 3.29 Subject vs. no-control costs: $\log_2$ difference - practice effects

The main cause for the observed variation seems to be the outcome of the no-control rule, rather than the outcome of the subjects' decision-making. Remember that exogenous sales are random and are different from trial to trial. Without the presence of any control at all and thus without the presence of correcting negative feedback, differences in exogenous sales are directly translatable into differences in costs. As exogenous sales vary from trial to trial, so do costs.

In contrast, random differences in exogenous sales should have a much smaller impact on sales in the presence of a control policy. Thus, we can assume that the control applied by the subjects smooths out the random differences in exogenous sales.

Subject effects

ANOVA Figure 3.27 shows that the subject effect was significant, ($p<.01$). Figure 3.30 gives an overview of subjects' performance. All subjects outperform the no-control rule in most of their trials; two subjects outperform the no-control rule in all of their trials (S12,S15).

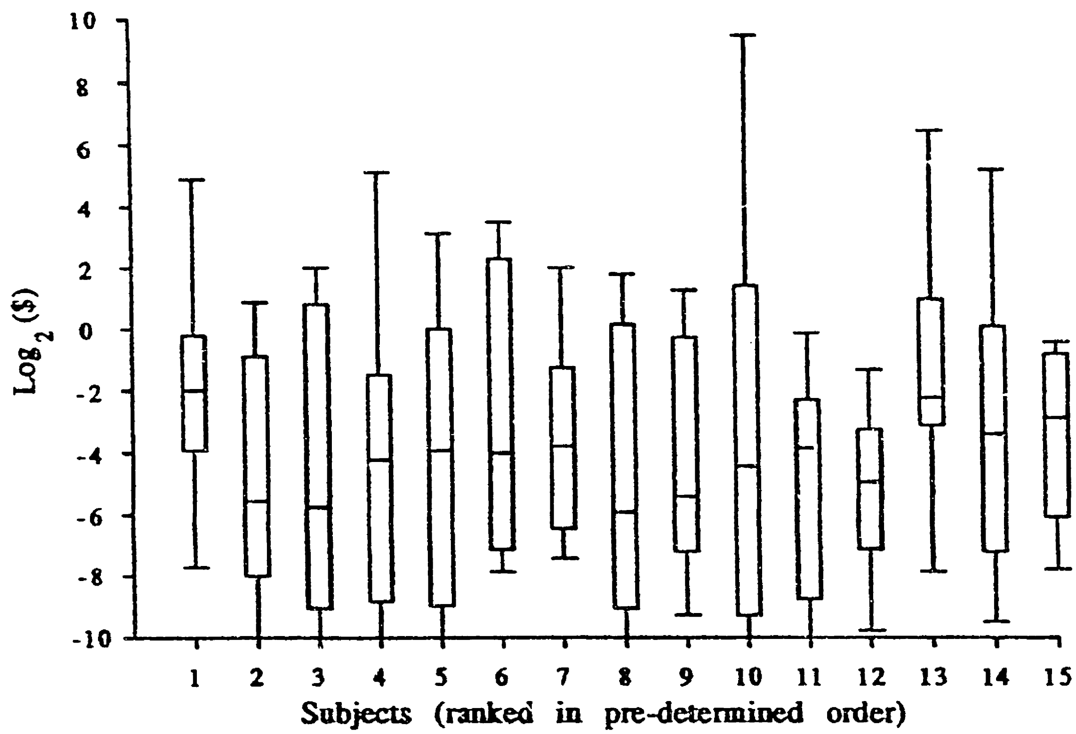

Figure 3.30 Subject vs. no-control costs: $\log_2$ difference -subject effects

Performance varies considerably between subjects. While the lowest ranked subjects (S1,S13) outperformed the no-control rule by 4 times at their median, the highest ranked subjects(S3,S8) outperformed the benchmark by more than 64 times at their median. While some subjects show a high consistency in outperforming the no-control rule (S11,S12) others show considerable variance (S4,S5,S6,S10).

3.3.6 Discussion of subjects' performance vs. benchmarks

Treatment effects

The results so far show a clear effect of delay and gain on the cost achieved. That effect is mainly due to the objective difficulty of the task. As the objective difficulty of the task increases, so do subjects' costs.

The question of which models subjects are using must await a detailed analysis of their notebooks and linear regression analysis. The convergence of the three analyses will provide a better picture of subjects' heuristics than one of the analytical methods alone. Hence, it seems premature to discuss the hypotheses based solely on the score data.

Practice effects

As discussed in section 3.2.3, on the no-control rule, the practice effect is attributable to two parts: learning and random walk effects.

Learning effects are difficult to assess because of the noise caused by the random-walk patterns of exogenous sales. It seems clear, however, that the first two trials show worse performance, and a plateau is reached from which we see only a slight increase in performance, if any, over the trials, showing that the initial training given to subjects prior to the experiment did have the desired effect of moving them quickly through the orientation stage.

Subject effects

Under increasing delay and gain, there seems to be a particularly strong deterioration in performance among those subjects who performed worst. Subjects who do poorly in the easier trials do especially badly in the most difficult trials as evidenced by the increased variance in the more difficult conditions. Hence, there seem to be differing plateaus for different skill levels. Rapoport and Ebert also struggled with the differences between subjects. For instance, Rapoport (1966b) modified his adaptive decision model in order to accommodate differences in the horizon with which subjects assessed the costs for various choices. His regression models suggested that subjects utilized various horizons for maximizing their payoffs.

As shown in Figure 3.30, these comparisons show substantial differences between different subjects which were expected; however, it will be one of the tasks in this thesis to illuminate why those differences arise. A more detailed discussion will await the notebook and regression analysis. The form of this analysis will assess the extent to which subjects utilize different frameworks, based on the cues that they consider and the weights given to them, and will combine the frameworks with different theories of how the frameworks exhibit dysfunctionality.

Various heuristic biases could also be responsible for the subject differences. Since the training did seem to have the effect of familiarizing the subjects with the task sufficiently such that little learning was needed with which to begin control, we can assume that the biases probably are not arising from an acquisition phase of information collection. In fact, this was an area that was controlled for in the design of the study. The biases are probably arising, then, from the processing of the information, which suggests that further analysis should focus on the manner in which people are combining the cues to formulate their decisions. The notebook analysis seems particularly well suited for determining, among those subjects that wrote calculations, what type of strategy they were employing.

3.4. Time data

We will use time spent on initial decisions and time spent on later decisions to distinguish between the competing hypotheses. Time spent on initial decisions is the average time per decision subjects spend in making their first decision in all games. Later time spent on decisions is the average time spent per round following the first decision of all games. In the discussion section, we will discuss the different hypotheses in light of the time data. Beyond the basic treatment, practice, and subject analyses, we will perform a regression analysis in order to determine the correlation between good and bad performance relative to the time spent per decision.

3.4.1. Time spent on later decisions

Overall performance

Figure 3.31 shows the mean time spent on later decisions by treatment conditions. The minimum time spent on a later decision was .22 minutes, and 1.6 minutes was spent as the longest time. The average time spent on later decisions was .54 minutes.

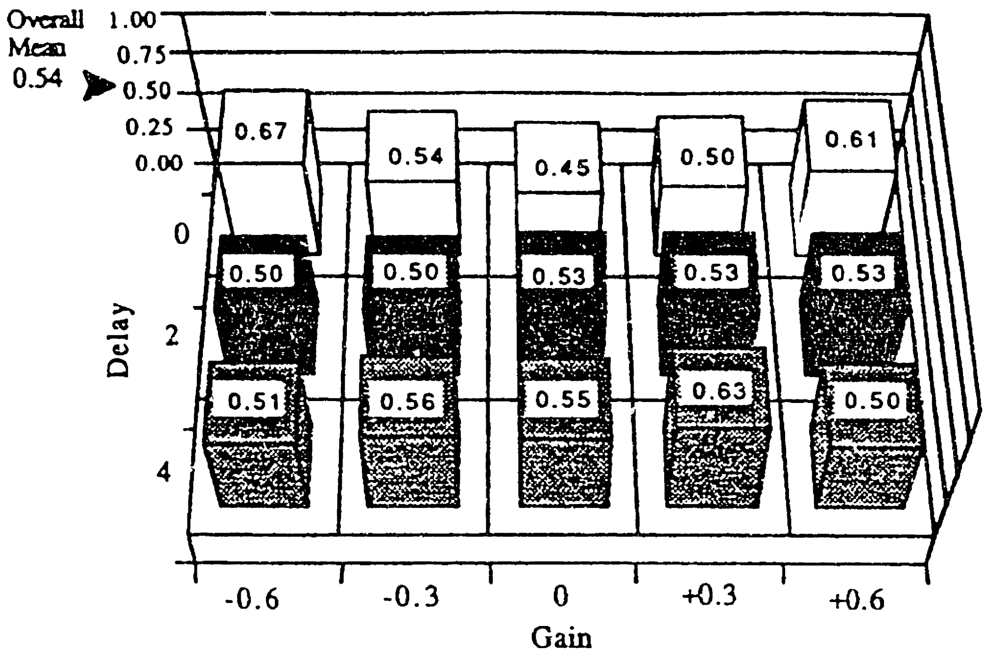

Figure 3.31 Average time spent on later decisions: minutes spent across treatments

Analysis

Source	df	F	probability
Overall treatment	14,182	1.88	p<.05
Delay	2,182	1.11	-----
Gain	4,182	0.70	-----
Delay x Gain	8,182	2.67	p<.01
Subject	14,182	15.90	p<.01
Practice	14,182	28.60	p<.01

-- not significant

Figure 3.32 Average time spent on later decisions: minutes spent- statistics

Figure 3.32 reveals a treatment effect attributable to the interaction of delay and gain (p<.01). Both subject and practice effects are significant (p<.01).

Treatment effects

From Figure 3.33, there does not appear to be large differences in the average time spent per condition; however, ANOVA Figure 3.32 reveals an overall treatment effect, attributable to the interaction of delay and gain, but it is not clear how to interpret this interaction. The minimum time spent on a later decision was .22 minutes, and 1.6 minutes were spent as the longest time. Subjects spent .54 minutes on average on later decisions.

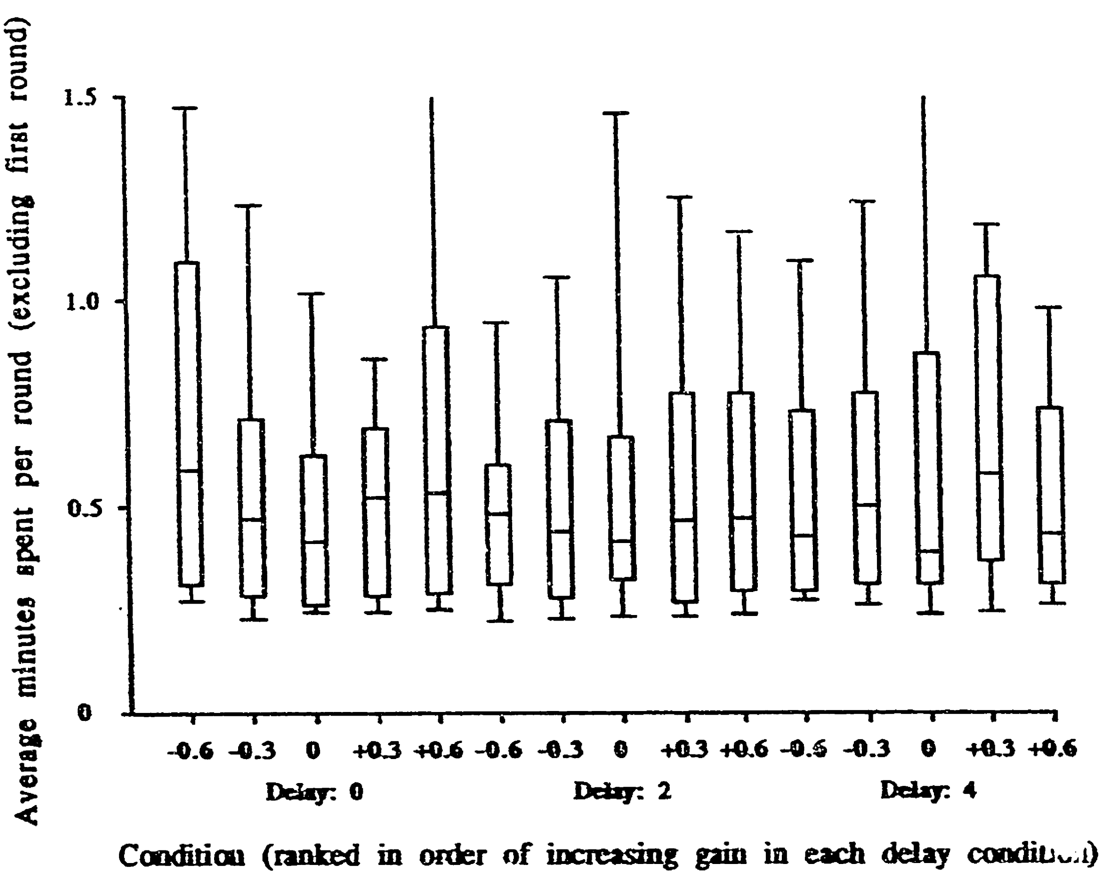

Figure 3.33 Average time spent on later decisions: minutes spent- treatment effects

Practice effects

As Figure 3.32 shows, significant differences between trials exist ($p<.01$). Figure 3.34 shows time spent on the later decisions for the 15 consecutive trials.

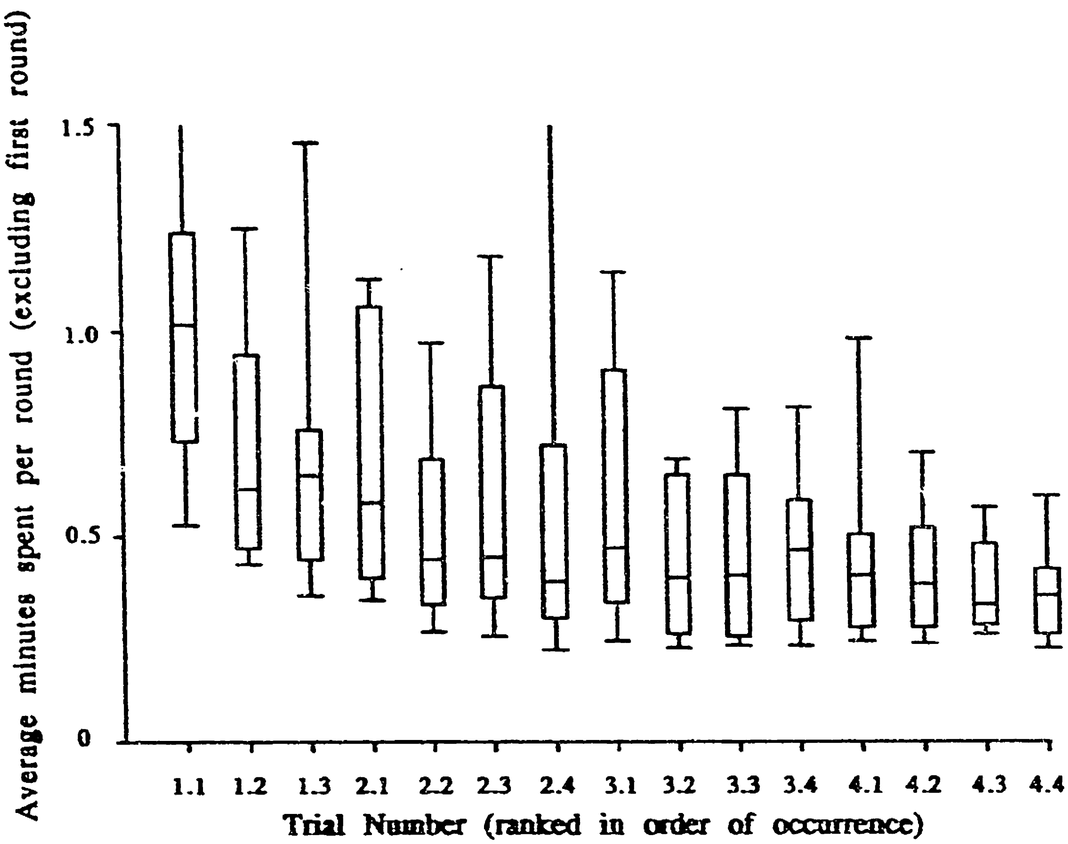

Figure 3.34 Average time spent on later decisions: minutes spent- practice effects

The median time, being slightly above 1 minute for the first trial, falls to .4 minutes for the second game in the second session (5th trial). In the fourth session it varies between .3 and .4 minutes. The minimum time falls from .5 minutes for the first trial to .23 minutes in the fifth trial and stays at this level for the remainder of the trials. The maximum time spent in each trial falls from 1.53 minutes for the first trial to .67 minutes for last two trials, with some spikes interrupting the broad pattern. The third longest time shows a similar pattern falling from 1.23 minutes in the first trial to .4 minutes in the last one.

In summary, a strong practice effect is observed during the first four trials. After that, the minimum time does fall further and the median time shortens only slightly. However, the spread between minimum and maximum time continues to narrow until the last trials, as the maximum time continues to decrease.

Subject effects

ANOVA Figure 3.32 shows significant differences between subjects (p<.01), and Figure 3.35 shows the time each subject spent on the later decisions of each trial.

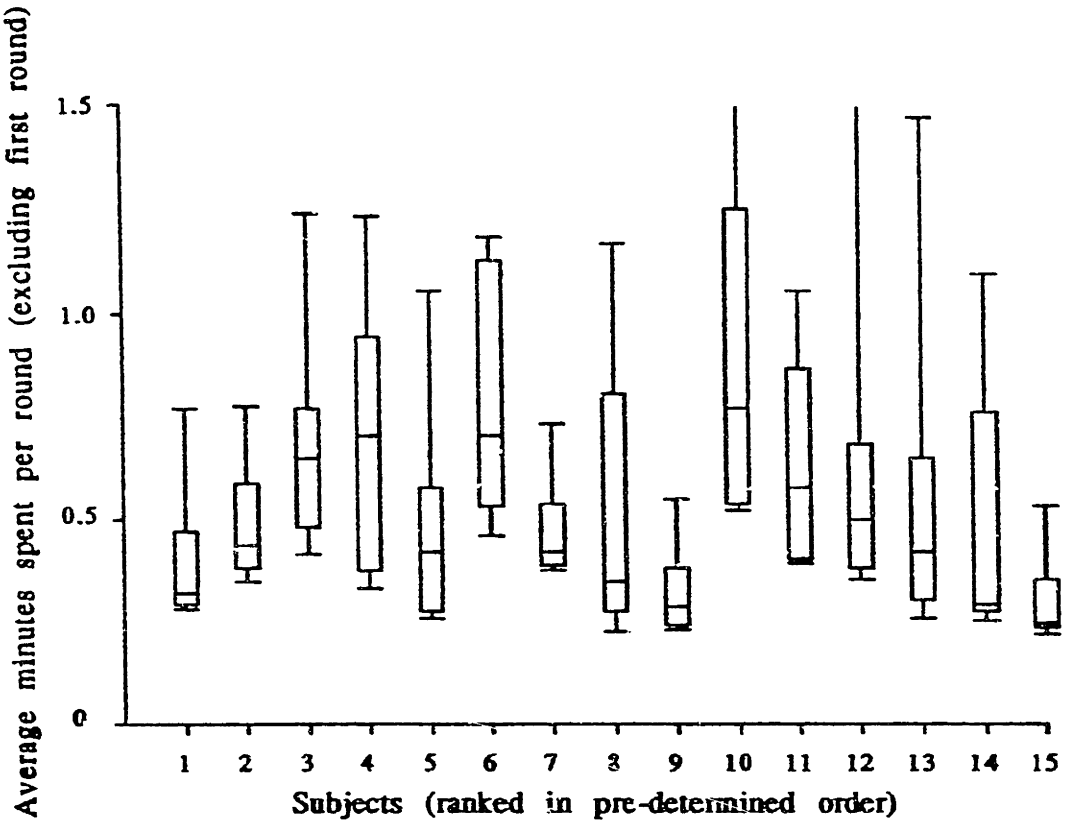

Figure 3.35 Average time spent on later decisions: minutes spent- subject effects

The subjects are shown in pre-determined order. Differences between subjects are apparent. The median time spent ranges from .27 minutes (subject 15) to .77 minutes (subject 10). The minimum time spent on a trial ranges from .23 minutes (subject 8,15) to .53 minutes (subject 10). While the difference between the 3rd shortest and the 3rd longest trial is below .17 minutes for some subjects (subject

1,7,9,15), it is more than .5 minutes for others (subject 4,6,8,10).

Regression Analysis

To determine the correlation of time spent and score achieved, subjects were ranked according to the $\log_2$ difference between subjects' costs and a combined benchmark of the average of optimum and no-control rules' costs. [3]

As time spent increases, we would have expected that score would have increased, but this assumption is not supported. See Figure 3.36, which shows that there are no clear differences between subject's performance attributable to effort in later decision time spent.

[3] The $\log_2$ differences ranged from -2.399 to +1.652. Refer to the notebook analysis, page 211, for the exact formula used to determine the $\log_2$ difference and a complete list of the ranked scores for all 15 subjects.

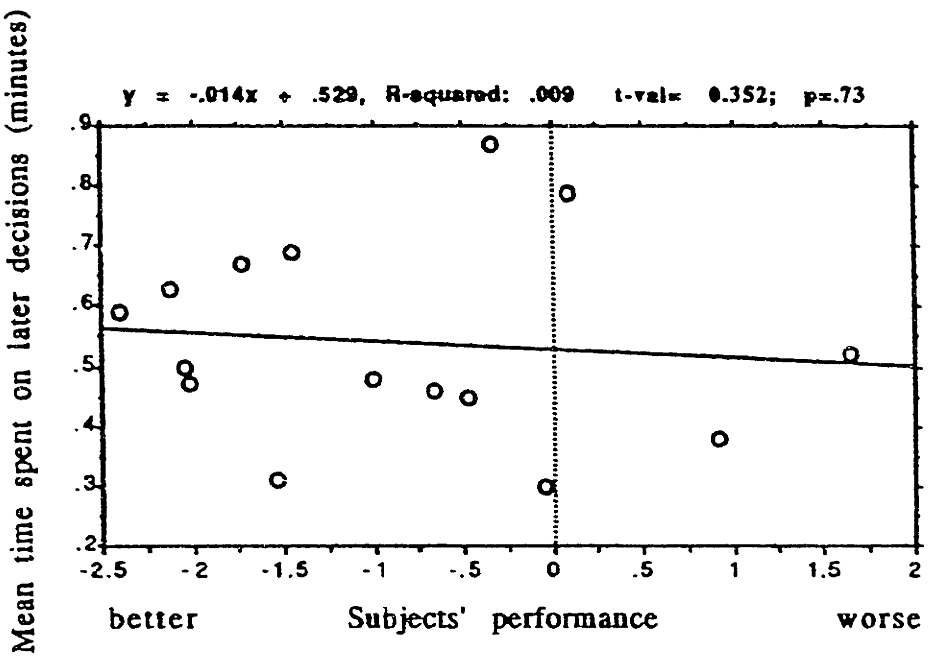

Figure 3.36 Average time spent on later decisions: regression against cost per subject

3.4.2. Time spent on the first decision

The time that is spent on the first decision in each game deserves to be treated separately. This data point could conceivably provide insight into the question of to what extent subjects develop a strategy at the beginning of each game that they apply throughout the rest of the game. It can be argued that the time spent on the first decision is spent on three separate activities: (1) Reviewing the new delay length and gain that will be in effect during the trial, (2) formulating a strategy in response to these conditions and (3) applying the strategy to compute the first decision.

Overall performance

Figure 3.37 shows the mean time spent on the first decision by treatments. The time spent on the first decision ranged from a minimum of .19 minutes to a high of 17.90 minutes. On average, 1.62 minutes were spent on the first decision. From the figure, we can see an increased time spent on the first decision in the 0 delay, -0.6 gain condition, which probably represents a familiarization effect. In addition, average time spent does not seem to vary considerably across conditions.

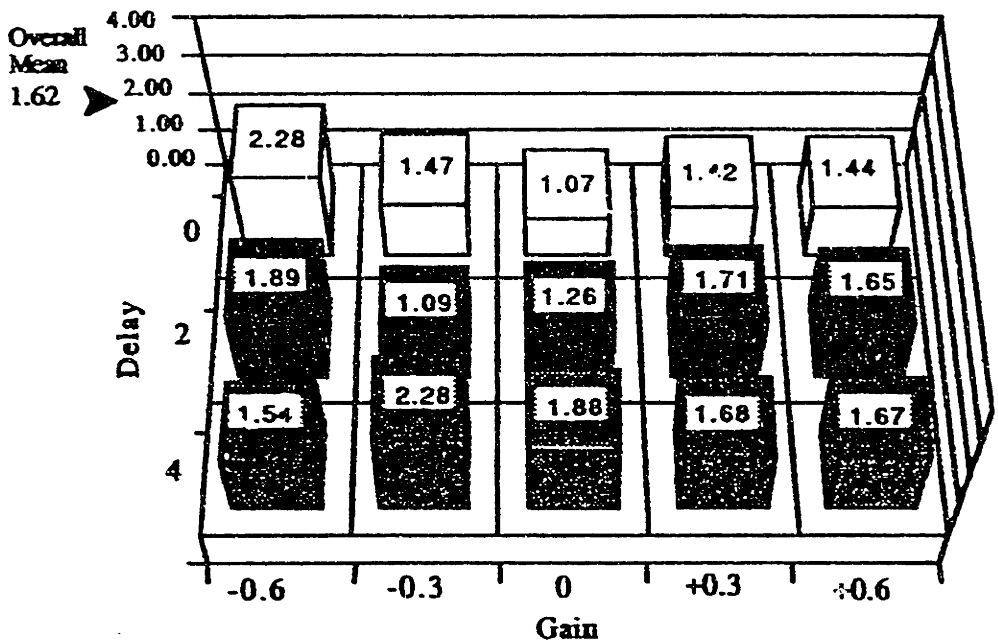

Figure 3.37 Time spent on first decision: minutes spent across treatments

Analysis

Source	df	F	probability
Overall treatment	14,182	0.93	-----
Delay	2,182	0.93	-----
Gain	4,182	0.68	-----
Delay x Gain	8,182	1.05	-----
Subject	14,182	3.02	$p<.01$
Practice	14,182	10.69	$p<.01$

-- not significant

Figure 3.38 Time spent on first decision: minutes spent- statistics

From ANOVA Figure 3.38, no treatment effects are present, but subject and practice effects are significant ($p<.01$).

Treatment effects

As ANOVA Figure 3.38 shows, neither a treatment effect of delay length nor of gain can be detected from the data. From Figure 3.39, subjects seem to take the same amount of time for their first decision no matter which condition they face.

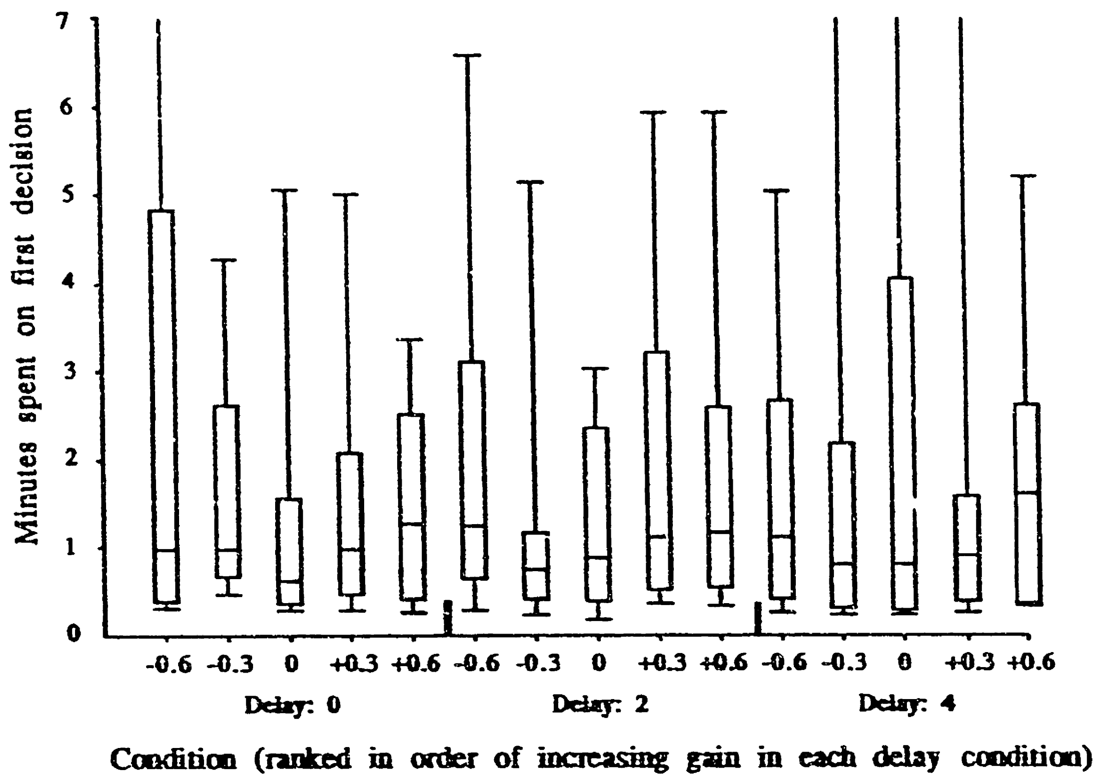

Figure 3.39 Time spent on first decision: minutes spent- treatment effects

Practice effects

ANOVA figure 3.38 shows significant differences between trials ($p<.01$). Figure 3.40 shows time spent for the first decision as a function of the trial sequence.

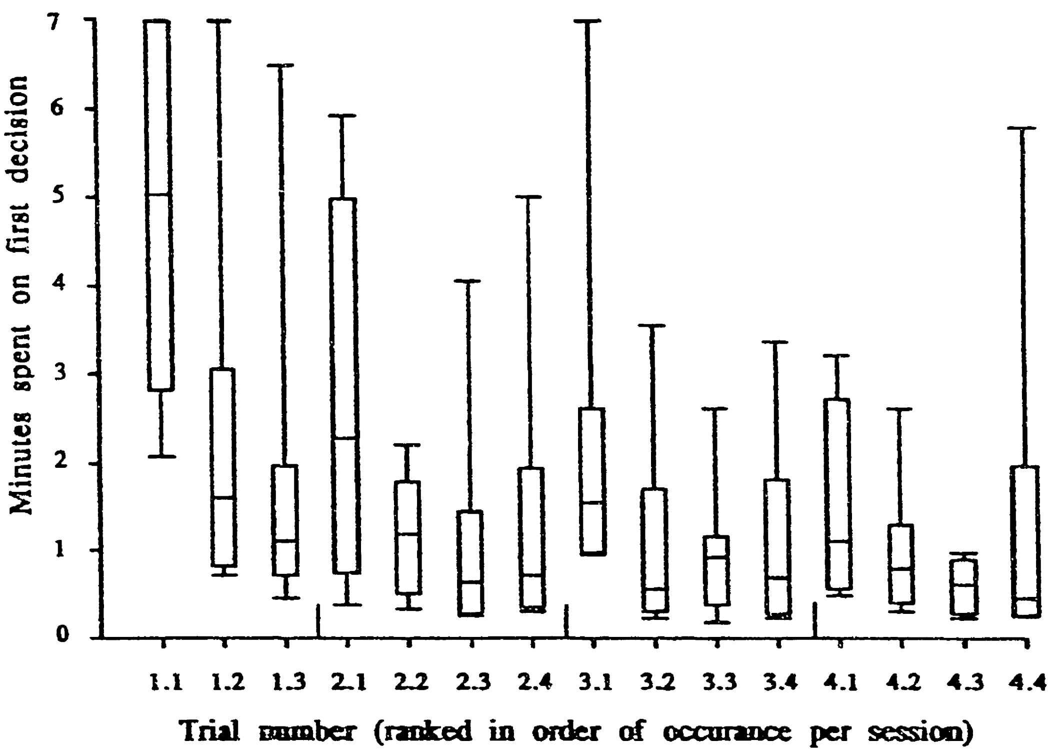

Figure 3.40 Time spent on first decision: minutes spent- practice effects

Median time spent on the first decision decreases from 5 minutes in the first trial to 40 seconds in 6th trial. It roughly remains at this level throughout the rest of the trials. A "beginning-of-session effect" overlays the general practice effect. Subjects spend more time at the beginning of the first trial in each session (1.1, 2.1, 3.1, 4.1) than at other trials. This effect is not unexpected, since subjects need some time to reorient themselves.

Subject effects

Figure 3.38 reveals significant differences between subjects ($p<.01$). Figure 3.41 shows subjects' differences. While some subjects spend a median time of only 20 seconds (S1,S15) on their first decision, the median for other subjects is almost six times higher (S4,S13).

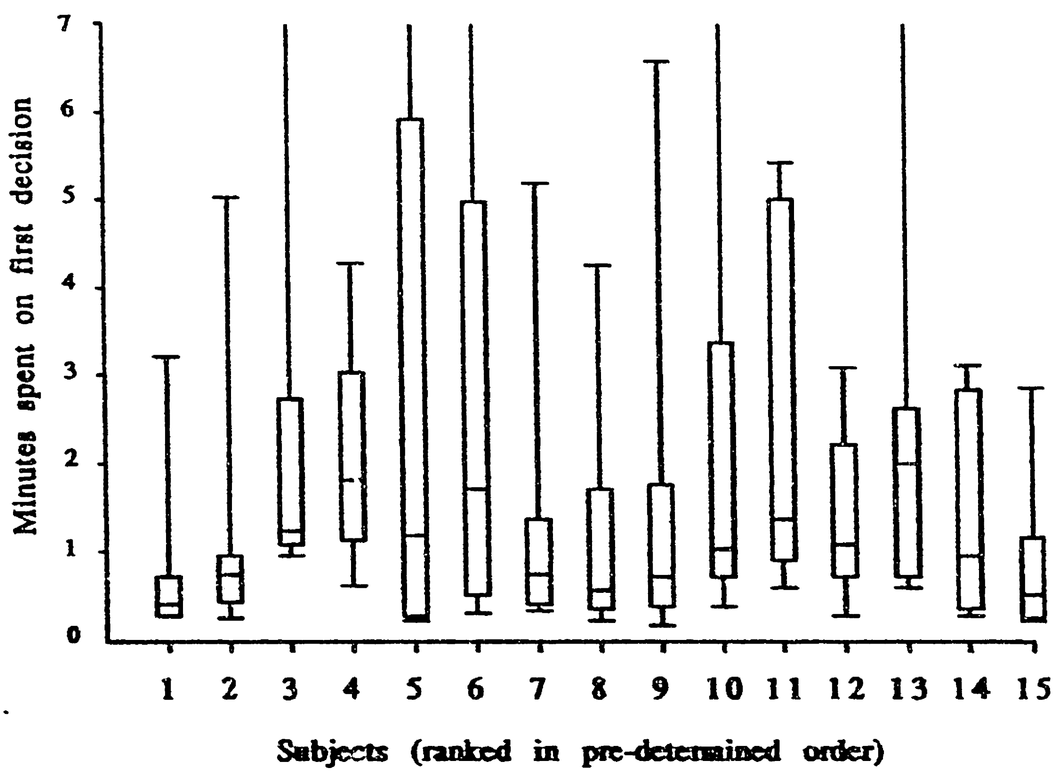

Figure 3.41 Time spent on first decision: minutes spent- subject effects

Some subjects show a variance of more than four minutes for the bulk of their first decisions (S5,S6,S11), while others show a variance of less than one minute (S1,S2). Again, the observed significant difference between subjects' behavior is not unexpected.

Regression Analysis

A regression analysis of time spent on first decisions against differences in scores achieved revealed no significant correlation between these two measures. As time increases, we would have expected that score would have increased; this assumption was not supported. See figure 3.42.

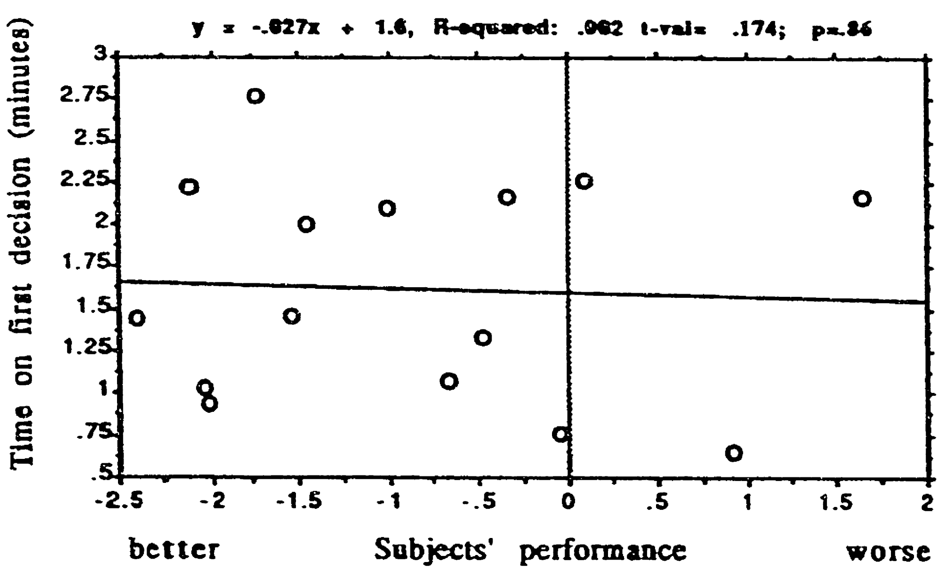

Figure 3.42 Time spent on first decision: regression against cost per subject

3.4.3 Discussion of Time Data

Treatment effects

As shown in Figure 3.37, the mean for the subjects' time spent on the first decisions is about one and a half minutes, while the mean time spent in the later decisions, from Figure 3.31, is approximately one half minute. Given the short decision times and given that subjects need some time to orient on each screen, the time data suggest that subjects cannot be performing very sophisticated calculations.

The lack of treatment effects speaks in favor of decision rules which do not take additional information into account with increasing delay and gain. Furthermore, these results support a mathematical model that takes the same amount of effort. Surprisingly, this is not the pattern of results that were expected. Score data suggest that more difficult decision situations, based on the objective difficulty found in the optimal rule simulation analysis, would require more time spent on a decision. In fact, we see that slightly less time is spent on the more difficult decision conditions.

An offsetting hypothesis, however, could be that as the difficulty of the task increases, the time spent on the decisions decreases; in essence, in easy conditions subjects spend more time because they can, with reasonable mental effort, "figure the system out", while in more difficult conditions subjects may simply give up. We cannot distinguish between either of these two views, nor can we rule out other possible explanations, until the more detailed notebook analysis.

Time spent on later decisions provides an opportunity to test the three central hypotheses of this thesis.

The _rational_ hypothesis would predict that time spent per trial increases with length of delay but is independent of gain. The complexity of the optimal solution and thus the amount of basic mathematical operations to be solved increases as the square of the number of states in the system. In contrast, the form of the solution is independent of the value of the matrix coefficients. (A value of 0 might be an exception, since it could simplify the overall solution.) In any case, a multiplication by 0.3 should be only slightly less or more difficult than a multiplication by 0.6. The data are not in accordance with the hypothesis: As delay increases, time spent per trial does not increase.

The _'Weight and Cue Adjustment'_ hypothesis is similar in its time data predictions to the rational hypothesis: As delay increases, the number of cues available increases. Thus, the time required to evaluate those cues and the time required to process them should increase. According to the same argumentation, strength of feedback should not have an impact on time spent since it does not affect the number of cues to be considered and to be processed. As in the rational hypothesis, the data are not in accordance with the hypothesis: As delay increases, time spent per trial does not increase.

The _'Weight Only Adjustment'_ hypothesis predicts that subjects concentrate on the most salient cues and ignore less salient ones. According to this hypothesis, subjects will ignore information on the additional states that are a result of an increase in delay time. The computational effort (and the time needed to execute these computations) will remain the same in all delay conditions. Similarly, gain will not change the computational effort involved in solving the task. The time data presented so far do not refute this hypothesis.

Time spent on first decisions also affords discussion of these hypotheses, as well.

The _rational_ hypothesis would require considerable computational effort at the beginning of each new trial. To derive the adequate decision rule for each condition, a quite complex optimization problem has to be solved. The complexity of the problem increases roughly quadratically with the number of states. With the possible exception of a gain of zero, gain does not affect the complexity of the decision rule. The rational hypothesis predicts long average reflection time at the beginning of each trial and predicts that the reflection time increases with delay length. Both predictions are not supported by the data.

Both the _'Weight Only Adjustment'_ and _'Weight and Cue Adjustment'_ have the same predictions as far as time spent at the beginning of each trial is concerned and can be discussed together. The hypothesis posits that at the beginning of each trial subjects adjust their overall decision rule in reaction to the new delay length and gain. While the hypothesis is not explicit about the mental effort required by the adjustment process, it can be assumed that the effort is far lower than what the rational hypothesis would predict. Assuming that a delay length of 0 and a gain of 0 would serve as the anchor, we would expect to find minimum reflection times under these conditions, since no adjustment would be required from the subjects. The data bear out this prediction to a limited degree only, if at all. Since neither a treatment effect of delay length nor of gain could be detected, a more detailed statistical analysis was not warranted. A visual inspection of the data reveals that the median of the 0 delay length, 0 gain condition is the lowest overall and the 0 gain conditions in general seem to be require less initial time than other gain conditions. Against expectations, initial time spent in the 0 delay length conditions was not lower than time spent in the longer delay length conditions.

In summary, the data do not contradict the heuristic hypotheses under the assumption that the mental effort required for the adjustment process is low compared with the mental effort required to compute decisions that follow from the rule derived.

Practice effects

As we would have expected from the literature, there is a strong practice effect for both later and initial decision times, shown in Figures 3.34 and 3.40. There seems to be a continuing time decrease until trial seven which is contrasted with the limited practice effect seen in the score data that decreased substantially for only the first two trials, after which performance levels off. While scores do not improve much after the first trials, time data are different, leading us to the belief that while subjects might not change their strategy after the first or second trial, they might increase the efficiency with which they apply whatever rule they use.

The time spent on the first decision which reflect practice effects is overlaid by a beginning of session effect, shown in Figure 3.40. The time spent on the first trial of each session is longer than the time spent on first decisions of trials later in the session. Beginning session times indicate an orientation phase, while later times reveal an application phase.

Regression analysis: time spent compared to score

A regression analysis was performed in order to determine the extent to which differences in time spent are correlated with differences in scores achieved. As time increases, we would have expected that scores would have increased, but this assumption is not supported. Increased time does not translate into better scores, nor is the reverse true, that better scores are attributed to those subjects that spent more time. Figures 3.36 and 3.42 show that there are no clear differences between subject's performance.

Although there are differences in effort spent, as evidenced by a decrease in variability in the later decision times, we cannot find a relation to subject's performance. This result is somewhat surprising.

Thus, we conclude that the time spent per decision is not sufficient to support an optimal strategy which requires substantial mathematical calculations. Since there were no treatment effects with respect to time spent as more information cues become available, the subjects' processing time is unaffected, suggesting that subjects spend the same computational effort across treatment effects, speaking in favor of single to two cue constant hypotheses, and speaking against both the rational and multiple cue hypotheses.

3.5. Control Effort Spent

A first strategy to analyze what causes the score results is to look at how much control effort is spent by the subjects. From a control effort point of view, underperformance can be caused by either exerting too much or too little effort. The cost structure of the task provides us with a ready measure for control effort spent. As explained in chapter 2, the overall costs are composed of cost associated with a change in production (control effort) and cost associated with deviations of inventory from its setpoint (goal discrepancy). By computing the ratio of control effort cost versus goal discrepancy cost relative to optimum for all 225 trials we get a measure of how aggressively subjects control the system.

As we did in section 3.3.4, we can compare the subjects' results to results suggested by the no-control rule and the optimal rule.

A comparison with the no-control rule is trivial for the purposes of this section. By definition, control costs associated with the no-control rule are 0. All costs observed result from goal discrepancies. Thus, as long as subjects perform any control at all, their relative control effort ratio is infinitely higher than the one computed according to the no-control rule.

The comparison with the optimal rule, by contrast, can provide meaningful insight. For the analysis below, the following ratio is computed for each trial and subject:

(Subjects' cost of control effort / subjects' cost of goal discrepancy) / (Optimal rule's cost of control effort / optimal rule's cost of goal discrepancy)

Overall performance

Control effort in Figure 3.43 shows the average for each of the 15 conditions. Values under 1 indicate undercontrol, and values above 1 indicate overcontrol. On average, the ratio (0.38) is about three times lower than 1, indicating distinct undercontrol. Undercontrol is prevalent in all 15 conditions, and undercontrol seems to be most prevalent in the presence of longer delays.

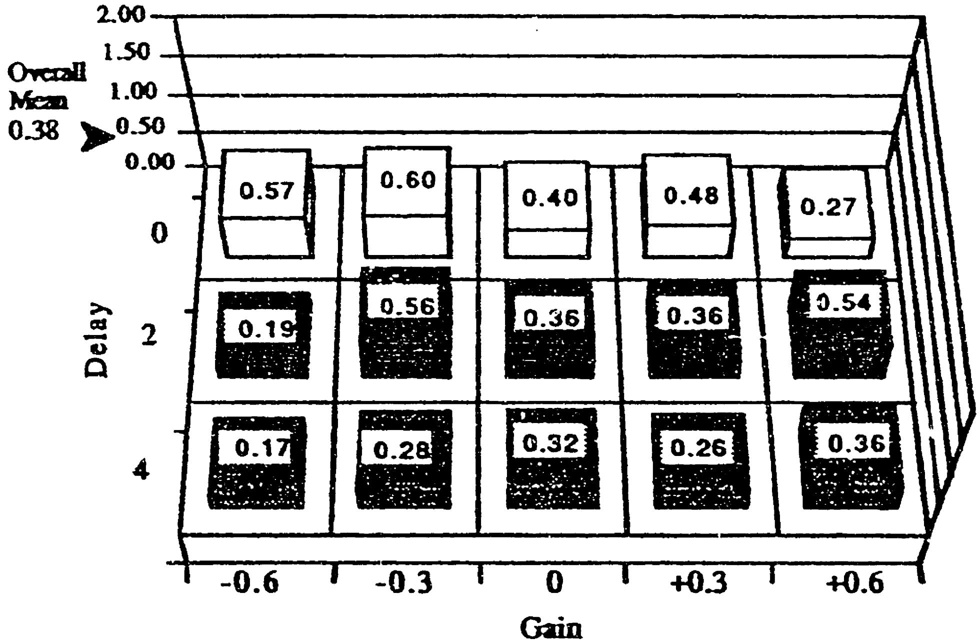

Figure 3.43 Ratio of subjects' vs. optimum control effort: ratio of cost across treatments

Analysis

Source	df	F	probability
Overall treatment	14,182	1.85	p<.05
Delay	2,182	4.20	p<.05
Gain	4,182	1.13	-----
Delay x Gain	8,182	1.63	-----
Subject	14,182	6.99	p<.01
Practice	14,182	1.00	-----

-- not significant

Figure 3.44 Ratio of subjects' vs. optimum control effort: ratio of cost- statistics

ANOVA Figure 3.44 reveals an overall treatment effect attributable to delay (p<.05), and significant differences between subjects (p<.01), while trials do not show differences in control effort.

Treatment effects

ANOVA Figure 3.44 shows a significant overall treatment effect (mostly attributable to the treatment effect of delay). Control effort Figure 3.45 shows the results in more detail. It appears that the median is highest in the "easy" conditions: Subjects control more aggressively in the absence of delays. In the presence of delays, subjects control more aggressively when there is no negative or positive feedback. It seems that subjects follow a common-sense heuristic: "Do not act, if you do not know what you are doing"

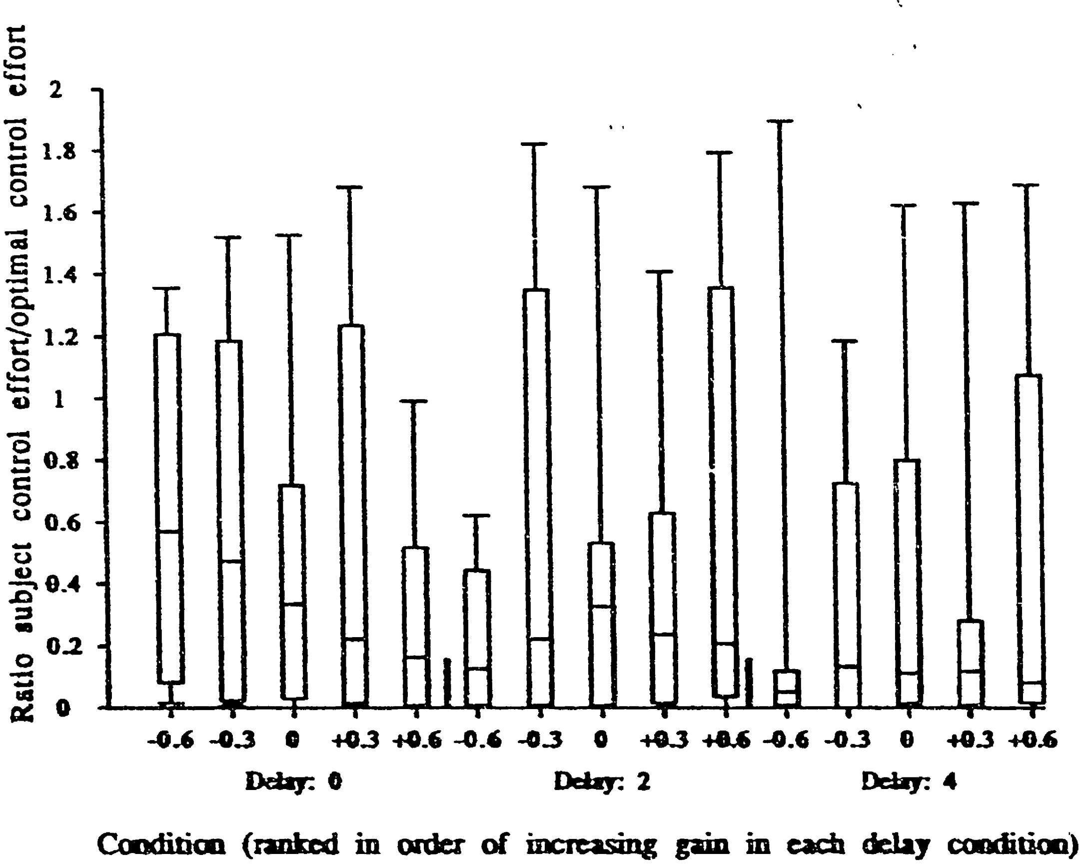

Figure 3.45 Ratio of subjects' vs. optimum control effort: ratio of cost- treatment effects

Practice effects

Figure 3.44 shows that practice effects are not significant. Control effort Figure 3.46 shows the control-ratio for the 15 consecutive trials. Undercontrol is most prevalent in the first trial, lending support to the hypothesis that undercontrol can be seen as a function of subjects' understanding of the task. Although a clear trend cannot be seen, it seems that control is higher in the latter trials than in the first five trials, lending further support to the hypothesis.

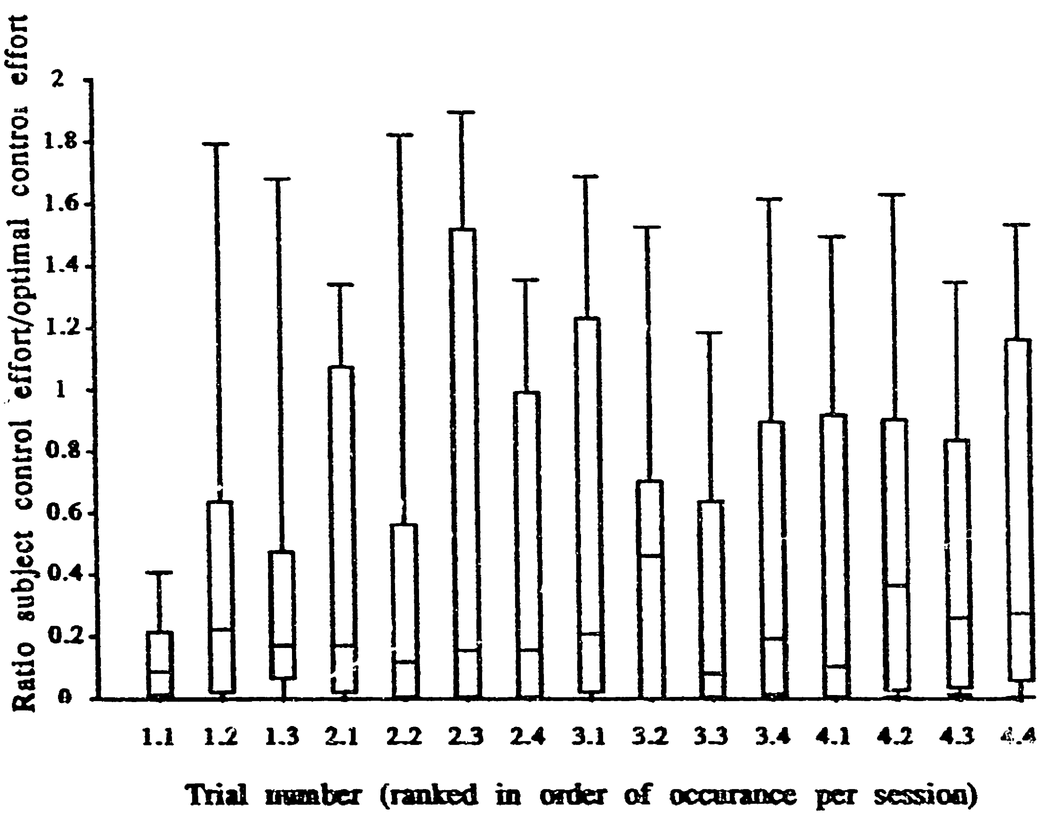

Figure 3.46 Ratio of subjects' vs. optimum control effort: ratio of cost- practice effects

Subject effects

Figure 3.44 shows that the subject effect is highly significant ($p<.01$). Control effort Figure 3.47 shows the control ratio for the 15 subjects in pre-determined order. Distinct differences are readily apparent: Four of the subjects (#1, #7, #13, #15) exert almost no control compared to the rest of their peers. According to the newly developed hypothesis, we would expect that these subjects are the low performers in the group. Regression analysis, chapter 5, will address this issue. With the exception of subject #8 and the four subjects mentioned above, all subjects overcontrol the system at least once.

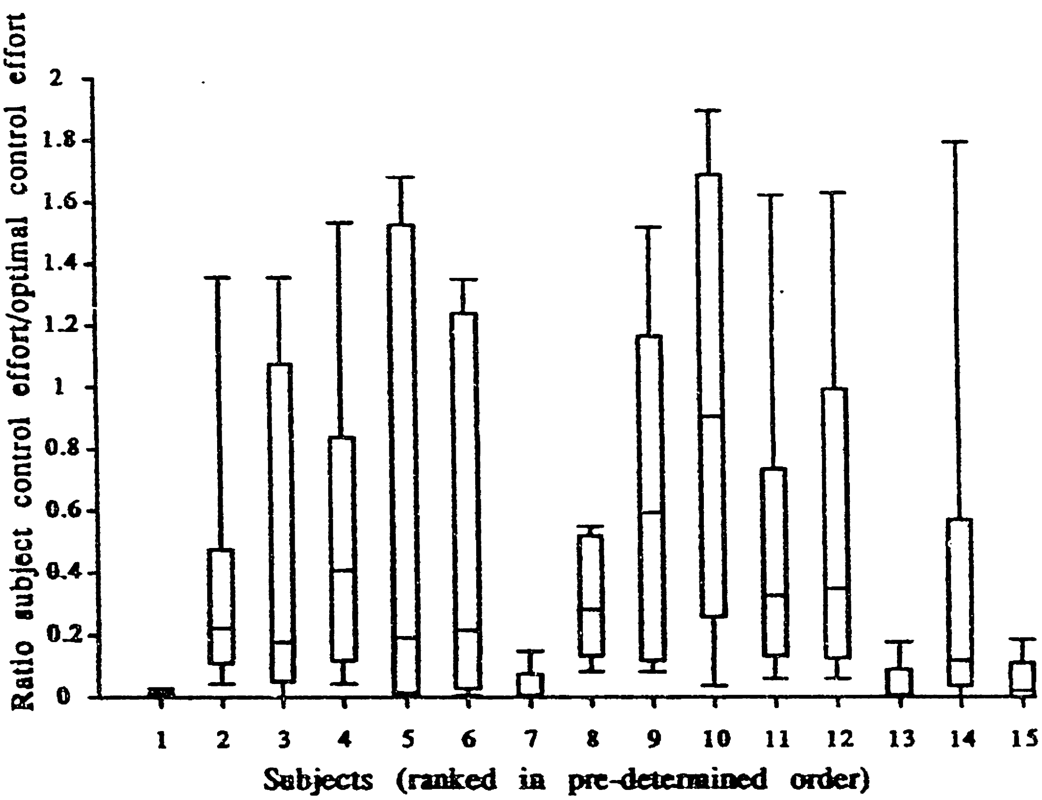

Figure 3.47 Ratio of subjects' vs. optimum control effort: ratio of cost- subject effects

Although we might fall into the trap of explaining randomness, the result could be seen as a sign that most subjects are willing to experiment with aggressive strategies at least once. The notebook analysis will help us shed more light on this tentative conclusion.

Regression Analysis

Subjects' performance was compared with the amount of control effort spent, as shown in Figure 3.48. We can see that there is definitely a relationship between score achieved and control effort spent ($p<.08$, $R^2 = .21$), indicating that greater control effort resulted in lower scores. If we treat subject 10 as an outlier, the relationship is much stronger ($p<.01$, $R^2 = .42$).

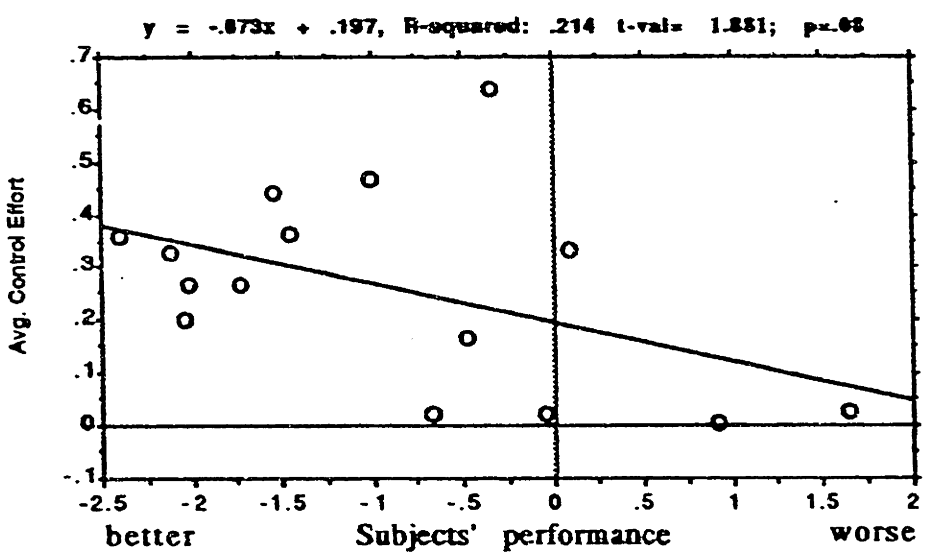

Figure 3.48 Ratio of subjects' vs. optimum control effort: regression against cost per subject

Discussion of Control Effort

The analysis of control effort gives us an explanation for subjects' underperformance that will be pursued further in the notebook and regression analyses. The very significant and surprising undercontrol throughout the conditions suggests an overly cautious approach, which worsens with increasing delay. The evidence suggests a clear refutation of the optimum rule hypothesis and a systematic bias in subjects' rules. It is somewhat puzzling, however, that despite the clear difference between subjects' and optimal control strategies, we do not see more of a difference in scores. We will revisit this paradox when we compare subjects' and optimal models and elaborate on them.

Figure 3.46 suggests that part of the learning that is achieved over the first two or three trials could be attributable to the fact that subjects gain greater confidence, albeit only slightly so. Further analysis suggests that this finding is tentative at best.

While none of the time data correlated with subjects' scores, undercontrol seems to be clearly correlated with scores, giving us a first insight into differences between low and high performers. We must await further analysis to determine the exact causes of undercontrol.

4.1 Introduction

Recall that subjects received a notebook in which they wrote their decisions in the upper right hand corner, while other calculations could be made on the rest of the page, thus providing some data about what calculations, if any, were made. The purpose of the notebook analysis, therefore, was to examine subjects' written calculations and comments in an effort to discern their strategies for performing the task. At the same time, the notebook analysis was intended to inform the regression analysis as described in chapter five. Extensive analysis of the notebooks consisted of reviewing subjects' round by round decisions, while relating the written calculations and comments to each decision. Since round by round analysis is a very time intensive task, we chose to perform the analysis on only one half of the subjects, selecting seven of them randomly.

The analysis was guided by the attempt to recreate the decision maker's problem space. In recreating the decision situation as it evolved over time, particular attention was given to all those situations where the different hypotheses that were introduced in section 2.7 and motivate the thesis lead to differences in their predictions. For example, in the presence of 1) a positive inventory gap,
2) overproduction relative to sales, and 3) large positive supply line all three hypotheses would predict a decrease in production. Yet, whenever some of the cues suggest increases in production while other ones suggest decreases, the hypotheses differ in their predictions of the decision maker's action.

Each subjects' data consists of 15 games of 32 rounds each. The analysis consisted of examining the initial games, 1 through 3, to see how subjects approached the task. The strategies of the initial games were compared with select games of sessions three and four with particular emphasis on long delay and positive gain, in order to determine how strategies formulated early on held up over time and under the difficult conditions. All games with long delays and high gains were investigated, since they reveal the most insight into the kind of understanding that subjects had achieved.

The analysis was guided by a behavioral check-list that had been constructed in advance. The list was motivated by previous research on misperceptions of feedback (Sterman, 1989a). Items on the list included:

1) average time overall and per game, including time differences within a game

2) many/few calculations

3) ignores gain

4) ignores delay

5) miscalculates gain

6) miscalculates delay

7) over/under emphasizes inventory

8) over/under emphasizes change in inventory

9) calculates gain appropriately

10) calculates delay appropriately

11) uses change in inventory appropriately

12) miscellaneous category: including an external/exogenous focus

In parallel with the notebook analysis, a more detailed examination of the time data was conducted. Time profiles for each game were analyzed in conjunction with notebooks to determine whether longer decision times, denoted by spikes in a time profile, related to changes in strategies. For example, subject 8's time profile for game 3 of session 1 shows (Figure 4.1) that he spent approximately .7 minutes on most decisions in that game; however, in rounds 10 and 28, the subject spent approximately 1.50 minutes per decision. Unexpectedly, these times did not correlate with any changes in strategy, as reflected in the decisions of the subject. It was necessary to check these types of effort outliers in order to understand each subjects' strategy for performing the task, since the increased times may have indicated changes in strategy.

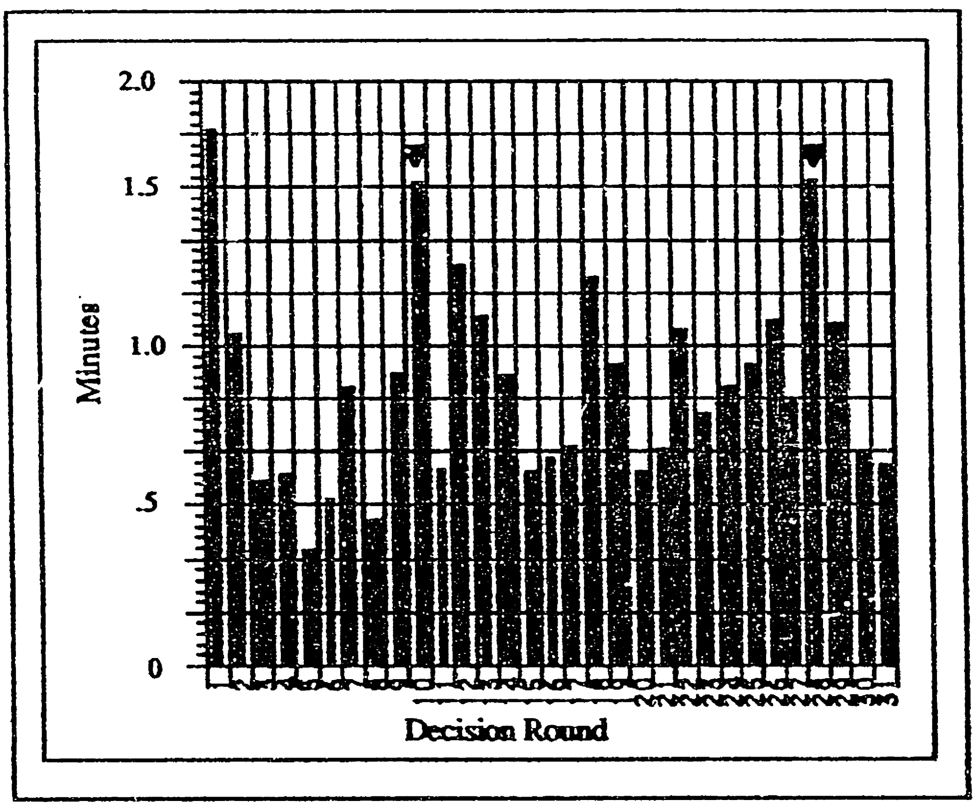

Figure 4.1 Example time profile

A note of caution seems appropriate. Despite the efforts made to code the notebooks objectively and to document accurately the extent to which the various cues are considered by the decision maker, the notebook analysis still might contain a certain amount of subjective inference. Throughout this analysis we try to guard against this tendency by supplying the readers with sufficiently detailed data about the decision situation to check the inferences we have drawn.

While we feel that the notebook analysis is a valuable addition to the researcher's toolkit of process methods (Carroll, and Johnson, 1990, pp.71-90), it has its inherent limitations. In particular, writing down calculations and performing them are two different processes. Just because nothing is in the notebook does not necessarily mean that subjects did not perform calculations in their heads.

Our analysis revealed three major levels of sophistication of control. Those three levels are:

1) attention to inventory only: subjects attempt to control the system by focusing solely on the inventory discrepancy.

2) attention to inventory and change in inventory: subjects understand the relationships between production and sales discrepancies and their importance.

3) attention given to inventory and expected change in inventory: subjects attempt to control the supply-line.

At the end of the following subsections, we will expand the summarized findings and place the different levels in the context of the prior hypotheses about control heuristics and our feedback framework.

4.2 Detailed analysis

Throughout the detailed analysis section we will discuss the subjects' decision situations with the help of a stylized display table, as shown in Figures 4.2 and 4.3.

round	60	61	62	63
change in production	0	?		
production	600	600	600	
dependent sales	-180			
independent sales	770			
total sales	590			
change in inventory	10			
inventory	10			

Delay = 2, Gain = -0.3
Subject 12, Trial 1.2

Figure 4.2 Stylized display table: Before decision in round 61

round	60	61	62	63
change in production	0	-11		
production	600	600	600	589
dependent sales	-180	-177		
independent sales	770	766		
total sales	590	589		
change in inventory	10	11		
inventory	10	21		

Delay = 2, Gain = -0.3
Subject 12, Trial 1.2

Figure 4.3 Stylized display table: After decision in round 61

We present the delay and gain conditions along with the subject and trial numbers at the bottom of each table. Within each table eight pieces of information are readily observable:

1) decision round
2) change in production
3) production and supply line
4) value for dependent sales
5) value for independent sales
6) total sales
7) change in inventory
8) inventory (surpluses or backlogs)

Refer to chapter 2, section 2.3 for a detailed description of the task.

1) In the stylized data tables, the decision round row shows the current decision round, denoted by "•" as well as any future decision rounds in games that contain delayed production. In Figure 4.3, which has a delay of 2, decision rounds are shown for the current round of 61 and the future rounds of 62 and 63.

2) A decision, change in production, has an immediate effect on dependent sales but may have a delayed effect on production. For instance, the subject in our example table made a decision to decrease production by -11 in round 61. Since this game has a delay of 2 the decision will not affect production until round 63 where production can be seen to drop to 589.

3) Values for production are displayed to the extent that they were visible to subjects on the actual game screen.

For instance, under a delay condition of two, subjects could view values for production two rounds ahead, as shown in figure 4.2. Throughout this section we will refer to future values of production either as the supply line or expected production.

4, 5, & 6) Dependent sales are determined by multiplying production with the gain factor. For example, in Figure 4.3 in round 61, dependent sales are determined by multiplying -0.3 by 589, which is the result of a change in production of -11 delayed two rounds. Independent sales are determined from the random-walk pattern for a given round. Total sales is the sum of dependent and independent sales. The change in inventory is the difference between current production and current total sales. The results of the subject's -11 production change can be seen in Figure 4.3.

7 & 8) Inventory, itself, is determined by adding the current change in inventory to the previous rounds' inventory. For example, 11 + 10 = 21, the value for inventory after round 61's production change, see Figure 4.3. Inventory is shown as either positive, negative or zero. A positive inventory implies a surplus in inventory, a negative inventory implies a backlog in inventory.

4.2.1 Subject 2

Early games

In game 1 (delay of 4 and gain of -0.6), subject 2 spends only about 5 rounds in an orientation phase in which he neither understands the calculation of gain nor the influence of delay. Quickly after the above start subject 2 develops a clearer understanding of the task. He notes the ranges, in round 66, of independent sales, " 969-999 range of sales," and he makes appropriate production changes that reflect an appreciation of all three information cues: inventory, change in inventory and supply line.

round	78	79*	80	81	82	83
change in production	0	10	0	0	0	0
production	575	575	585	585	585	595
dependent sales	-351	-357	-357	-357	-357	-357
independent sales	952	940	938	952	943	945
total sales	601	583	581	595	586	588
change in inventory	-26	-8	4	-10	-1	7
inventory	-2	-10	-6	-16	-17	-10

* = current decision round

Delay = 4, Gain = -0.6
Subject 2, Trial 1.1

For instance, at decision round 79, he saw from round 78 that inventory was -2 and that the change in inventory was -26. In the supply line, he determined that the difference between production and sales for the upcoming rounds could possibly become negative, further swelling the backlog. So, he writes, " I see problems ahead of low production." Thus, he increases production slightly which indicates consideration of the supply line. The next round reveals an increase in the backlog but a decrease in the change in inventory. As a consequence, he decides not to increase production further, indicating an appreciation of the future change in inventory.

Although subject 2 learns to control the system fairly well in the first game, in round 85 he shows some tendency to rely, unrealistically, on exogenous factors to come to his aid. There is some indication that the subject does not only take independent sales into account, but also that he hopes and relies in his strategy that independent sales will come to his aid.

round	84	85	86	87	88	89
change in production	0	0	-10	-10	-25	0
production	595	595	595	595	595	595
dependent sales	-357	-357	-351	-345	-330	-330
independent sales	**937**	**927**	**913**	911	923	932
total sales	580	570	562	566	593	602
change in inventory	15	**25**	33	29	2	-7
inventory	5	**30**	63	92	94	87

Delay = 4, Gain = -0.6
Subject 2, Trial 1.1

For example he writes in round 85, " Hope that independent sales will go up." The subject states that he hopes independent sales will increase in order for the difference between production and total sales to decrease which would, in turn, stabilize or decrease inventory. His production cutback is very cautious, thereby indicating his hope that the environment will come to his aid, so that he does not have to incur the high cost associated with changes in production. The environment, however, disappoints him; independent sales drop from 927 to 913, and the subject states, " uh oh! Sales went down."

Games two (delay of 0 and gain of +0.3) and three (delay of 2 and gain of -0.3) reveal that subject 2 carries on with strategies 2 and 3, incorporating change in inventory and, to a lesser extent, expected change in inventory.

round	76	77	78	79	80	81	82
change in production	**-10**	**-10**	**-5**	**0**	**20**	**20**	**10**
production	540	530	520	510	505	505	525
dependent sales	-156	-153	-152	-152	-158	-164	-167
independent sales	690	698	706	711	698	707	699
total sales	534	545	555	560	541	544	533
change in inventory	6	-15	-35	-50	-36	-39	-8
inventory	120	105	70	21	-15	-54	-61

Delay = 2, Gain = -0.3
Subject 2, Trial 1.3

For example, in round 72 through 78 of game 3, the subject decreases production in response to a positive increase in inventory; however, his decreases continue too long causing a larger negative oscillation than necessary. Thus, he exhibits some tendencies to "ignore the supply line" (Sterman, 1989a).

Later games

By game 3 of session 4 (delay of 4 and gain of +0.3), subject 2 shows signs of supply line control, stage 3 strategy, either by attempting to compensate for oscillations in inventory before they occur or by attempting to correct perceived past mistakes.

round	71	72	73	74	75	76
change in production	0	-20	-30	-10	**20**	**20**
production	525	525	575	580	580	560
dependent sales	174	168	159	156	162	168
independent sales	391	397	400	414	412	397
total sales	565	565	559	570	574	565
change in inventory	-40	-40	16	10	6	-5
inventory	92	52	68	**78**	84	79

Delay = 4, Gain = +0.3
Subject 2, Trial 4.3

For instance, after having cut production in response to increases in inventory in rounds 72-74, subject 2 increases production for two rounds in order to keep a backlog from occurring. Even though the magnitude of his decision was too aggressive, leaving a positive inventory discrepancy of approximately 80, the behavior showed that he was attempting to control for previous production decisions.

Effort

In general, subject two does not write many calculations throughout the experiment. Regarding time spent on decisions, subject 2 generally spent more time on later decisions than on first decisions. On average, he spent approximately .50 minutes per choice. His decision times per game are punctuated by several spikes in later rounds, lasting approximately 1 minute, that may have indicated re-orientation phases or strategy evaluation followed by plateaus of decision times, indicating application phases of strategy.

4.2.2 Subject 4

Early games

In game one (delay of 0 and gain of -0.3), subject 4 starts off with stage 2 focus, revealing a brief orientation phase.

round	64	65	66
change in production	-20	0	10
production	600	600	610
dependent sales	-180	-183	-183
independent sales	791	804	790
total sales	611	624	607
change in inventory	-11	-24	3
inventory	17	-7	-4

Delay = 0, Gain = -0.3
Subject 4, Trial 1.1

Not only does the subject increase production in response to a backlog, but considers change in inventory as well. For example, after inventory is moved into a surplus of +17 in round 64, and the change in inventory is -11, the subject makes a production change in round 65 of zero, indicating that he understands that the next round's change in inventory will decrease the surplus. Furthermore, in round 66, he can see from the previous round that inventory is -7 and change in inventory is -24 both of which indicate a larger production change than merely, say, +7. He increases production by +10, a value larger than inventory gap alone, possibly indicating an influence of change in inventory. Thus, subject 4 exhibits knowledge of the relationships between production and sales and their influence upon inventory, i.e. stage 1 and 2 strategies.

In game two (delay of 2 and gain of -0.6), subject 4 considers the supply line and makes "what if" calculations to obtain inventory values for two rounds in advance. For example,

round	72	73•	74	75	76	77
change in production	0	-10	-10	-15	-5	-10
production	575	**590**	**590**	580	570	555
dependent sales	-354	-348	-342	-333	-330	-324
independent sales	919	918	909	908	897	889
total sales	**565**	570	567	575	567	565
change in inventory	10	20	23	5	3	-10
inventory	-11	9	32	37	40	30

• = current decision round

Delay = 2, Gain = -0.6
Subject 4, Trial 1.2

In determining the production decision for round 73, he writes:

590 - 565 = 25 => (implies Inv. of) 14

590 - 565 = 25 => 39 Lower

His calculations indicate that he is incorporating the supply line, albeit incorrectly, by subtracting the total sales of the previous round from the production of the current round and next rounds. He decides to lower production because he predicts that inventory will be 39 in two rounds.

In a later decision, he continues the same strategy. For example,

round	79	80*	81	82
change in production	-10	15	5	0
production	540	**545**	**535**	550
dependent sales	-321	-330	-333	-333
independent sales	890	881	879	877
total sales	**569**	551	546	544
change in inventory	-29	-6	-11	6
inventory	2	-4	-15	-9

* = current decision round

Delay = 2, Gain = -0.6
Subject 4, Trial 1.2

For decision round 80, he writes:

545 - 569 = -24 + 2 => -22

535 - 569 = -34 - 22 => -56

Although subject 4 is being very specific about determining what inventory will be two rounds in the future if the current conditions remain the same, his changes in production are not directly related to his calculations. Because he does not make "what if," predictive, calculations based on production changes of particular values, but rather makes predictive calculations based on constant sales and other factors, he acts as if his own decisions do not have an influence on sales. His predictive calculations do not contain the idea that dependent sales are determined by his production changes; therefore, his strategy is always reactive.

In game 3 (delay of 2 and gain of +0.6), subject 4 utilizes the same strategy. In a negative gain situation the previous strategy was adequate; however, in this game, his miscalculations of dependent sales and delay cause large oscillations and substantial cost, since the gain is self-reinforcing. For instance,

round	76	77*	78	79
change in production	200	500	500	500
production	90	90	290	790
dependent sales	174	474	774	1074
independent sales	150	158	166	171
total sales	324	632	940	1245
change in inventory	-234	-542	-650	-455
inventory	-516	-1058	-1708	-2163

* = current decision round

Delay = 2, Gain = +0.6
Subject 4, Trial 1.3

He expresses his surprise that his strategy is not working by writing in round 77, "I can't believe this is happening!" At this point in the game, when backlog is increasing drastically, the subject abandons the previous strategy and adopts strategy 1. He ignores change in inventory as evidenced by large oscillations in inventory. In addition, he's no longer writing down any calculations, and his decision times are averaging approximately .5 minutes, thereby indicating a relatively quick, reactive strategy.

In session two, game two (delay of 4 and gain of +0.3), subject 4 exhibits a similar delay-gain strategy to the above, except that for approximately one-third of the game, he calculates current production decisions as if the production changes were based solely on the production value delayed four rounds, while he continues to use the total sales from the round previous to the current round. The early strategy is shown below:

round	70	71*	72	73	74	75
change in production	-7	**-14**	-18	5	-10	-5
production	602	590	601	609	**602**	588
dependent sales	181	176	171	173	170	168
independent sales	407	396	404	392	391	391
total sales	**588**	572	575	565	561	559
change in inventory	14	18	26	45	42	29
inventory	8	25	51	96	137	166

* = current decision round

Delay = 4, Gain = +0.3
Subject 4, Trial 2.2

The subject only writes:

602 - 588 = +14

He then makes a change in production of the same magnitude but opposite in sign, as illustrated above. By round 76, inventory is showing a large surplus of 187, at which time he decides to abandon his strategy and begins making larger decreases in production in -25 increments.

This later strategy is illustrated below:

round	75	76	77	78	79	80
change in production	-5	-5	-25	-25	0	20
production	588	570	575	565	560	555
dependent sales	168	167	159	152	152	157
independent sales	391	383	374	388	386	393
total sales	559	550	533	540	538	551
change in inventory	29	21	42	26	23	5
inventory	166	187	229	254	277	281

Delay = 4, Gain = +0.3
Subject 4, Trial 2.2

Before his changes have time to take effect, he begins to increase production again, as if he were trying to avoid the large oscillations seen in previous games; see decision in round 80. Throughout the rest of the game, inventory remains in a surplus state of approximately 60 units, which he does not rigorously attempt to control.

Later games

Throughout the later games, subject 4 ignores the supply line, which causes him great difficulties. The last rounds of the subject's last trial provide us with a unique opportunity to summarize the level of understanding the subject has attained.

round	86	87*	88	89	90	91
change in production	-500	**-5000**	**1000**	**1000**	**1000**	**1000**
production	700	200	0	0	-500	-5500
dependent sales	-300	-3300	-2700	-2100	-1500	-900
independent sales	191	206	200	194	185	170
total sales	-109	-3094	-2500	-1906	-1315	-730
change in inventory	809	3294	2500	1906	815	-4770
inventory	**254**	3548	6048	7954	8769	3999

* = current decision round

Delay = 4, Gain = +0.6
Subject 4, Trial 4.4

In round 87, he enters -5000 intending to enter -500, which would have been equivalent to his previous decrease in production. Up to this point, the subject has learned that he should be cautious, since he gets into trouble with large production changes, and he has gained minimal understanding of the importance of future change in inventory. Yet, he feels so uncomfortable with the system that he suspends control until the system is returned to a state in which he does feel comfortable. Again, he "suspends reality" in a sense, by attempting to make five consecutive increases in production of +1000 in order to make up for the typing error of -5000 before he can initiate control upon the system again. Thus, he applies his "insights" in the wrong sense. In this particular situation, it would have been much better to have compensated for the -5000 error right away and then resume control again.

Effort

In terms of time spent on decisions, subject 4 spends on average approximately .72 minutes per decision. His decision times are punctuated by spikes of longer times followed by quicker decision times.

4.2.3 Subject 8

Early games

In game one (delay of 2 and gain of +0.6), subject 8 exhibits a weak strategy 2, and weaker appreciation for supply line, strategy 3. As noted above, his performance is characterized by cautiousness and weak attention to change in inventory. His cautiousness takes the form of a "act and wait" strategy, i.e. reactive to inventory. For example:

round	65	66	67	68	69*
change in production	0	5	5	5	**150**
production	615	615	615	620	625
dependent sales	369	372	375	378	468
independent sales	264	250	252	260	273
total sales	633	622	627	638	741
change in inventory	**-18**	**-7**	**-12**	**-18**	-116
inventory	-69	-76	-88	**-106**	-222

* = current decision round

Delay = 2, Gain = +0.6
Subject 8, Trial 1.1

Prior to round 69's increase in production, the subject has practically ignored the change in inventory, as evidenced by the substantial backlog, the negative change in inventory, and the magnitude of his decisions. In round 69 he attempts to control more aggressively, but his later decisions reflect that he merely continues with the cautious bias and an "act and wait" strategy.

In rounds 84 through 87 of game one, for example, subject 8 makes small production changes in response to large surplus inventory and waits to see the effects of [illegible]is changes before initiating further decreases.

round	82	83	84	85	86	87
change in production	0	-40	0	0	0	-20
production	600	500	500	460	460	460
dependent sales	300	276	276	276	276	264
independent sales	223	225	217	207	193	191
total sales	523	501	493	483	469	455
change in inventory	77	-1	7	-23	-9	5
inventory	427	426	433	410	401	406

Delay = 2, Gain = +0.6
Subject 8, Trial 1.1

Immediately following round 87, the subject performs two larger increases in production in order to compensate for any oscillations caused by his previous decreases in production.

For instance,

round	88	89	90	91
change in production	**80**	**80**	**-20**	-40
production	460	440	520	600
dependent sales	312	360	348	324
independent sales	203	212	213	214
total sales	515	572	561	538
change in inventory	**-55**	-132	**-41**	62
inventory	351	219	**178**	240

Delay = 2, Gain = +0.6
Subject 8, Trial 1.1

Thus, he shows greater appreciation for inventory control, and the supply line, but his emphasis on change in inventory is weak. He does, in round 90, show a quicker effort to control for previous decisions in that his control for oscillations comes immediately after the increases in inventory instead of waiting for several rounds.

In game two (delay of 0 and gain of +0.6), subject 8 again graphically keeps track of production, inventory, and cost of inventory. His control strategy overall is level 2, and he maintains a cautious approach to decreasing surplus or backlogs. In rounds 69 through 72, he starts to show a stronger appreciation for change in inventory, seen below.

round	69	70	71	72
change in production	-10	**0**	**10**	**5**
production	490	490	500	505
dependent sales	294	294	300	303
independent sales	203	212	203	199
total sales	497	506	503	502
change in inventory	**-7**	**-16**	**-3**	3
inventory	**29**	**13**	**10**	13

Delay = 0, Gain = +0.6
Subject 8, Trial 1.2

In round 70, he sees that change in inventory in the previous round was -7 and inventory is 29, thus he makes no production change so that inventory will decrease. In round 71, he notes that change in production is larger than inventory and in the next round likely to cause a backlog, so he makes another increase in production of 10.

In game three (delay of 0 and gain of 0), subject 8 continues to make graphs of production, sales, and cost of inventory. He continues to exhibit strategy 2 and, similar to the previous game, he utilizes change in inventory effectively.

Later games

In game 3 of session 4 (delay of 4 and gain of +0.3), subject 8 has stopped making the previous graphs. He exhibits strategy 2 and maintains the cautious "act and wait" control style. In rounds 65 through 68, the subject shows the continuing non-aggressive control.

round	64	65	66	67	68	69
change in production	20	**0**	**0**	**0**	**0**	-20
production	600	615	645	685	705	705
dependent sales	212	212	212	212	212	205
independent sales	450	440	449	464	452	437
total sales	662	652	661	676	664	643
change in inventory	-62	-37	-16	**10**	**42**	63
inventory	-189	-226	**-241**	**-232**	**-190**	**-128**

Delay = 4, Gain = +0.3
Subject 8, Trial 4.3

In rounds 61 through 64, the subject has increased production by 105 in response to an increasing backlog. In rounds 65 through 68, the subject waits for his previous production increase to take effect.

Throughout the rest of this game, with small exceptions, the subject maintains the same pattern of waiting for his production changes to take effect before he takes actions. From round 71 until the end of the game, inventory remains in a large surplus; his production changes never are aggressive enough to decrease inventory. The subject displays strategies 2 and a weaker 3 and maintains the non-aggressive control posture that causes larger costs in inventory:

round	87	88	89	90	91
change in production	**20**	**0**	**0**	**0**	50
p.od. 'on	530	530	530	530	550
dependent sales	165	165	165	165	180
independent sales	393	393	387	373	372
total sales	558	558	552	538	552
change in inventory	**-28**	**-28**	**-22**	**-8**	-2
inventory	**223**	**195**	**173**	**165**	**163**

Delay = 4, Gain = +0.3
Subject 8, Trial 4.3

Effort

Subject 8 notes few calculations in writing, only graphically keeping track of production, inventory, and cost of inventory. His time profiles do not correlate w..h any peculiarities within his decision rounds. Overall, subject 8 spends approximately 1.25 minutes per choice. His later spikes are approximately 2 to 3 minutes long and did not correlate with any changes in strategy as evidenced by consistent decision behavior.

4.2.4 Subject 10

Early games

In game one (delay of 4 and gain of 0), subject 10 makes virtually no calculations and only writes costs of change in production and inventory throughout. She ignores the supply line and change in inventory, revealing a reactive control style, which leads to large inventory build ups. For instance, she begins the game with a small backlog of -3 which increases to -41 by the fourth round, meanwhile the subject has been increasing production up to 58.

At this point, she has put enough in the supply line to cover the backlog; however, she waits only one round, e.g. change in production = 0, and begins to increase production again. Like a person who has forgotten what speed she is already going and pushing the accelerator even further in order to keep up with traffic, subject two appears to want her production changes to take effect immediately, and when they do not increase inventory quickly enough, she increases production even further, causing later oscillations in inventory in the opposite direction, similar to results observed in the "Beer game" experiment (Sterman, 1989b).

round	63	64	65	66	67
change in production	**25**	**30**	0	**20**	**18**
production	600	600	603	613	638
dependent sales	0	0	0	0	0
independent sales	600	611	624	610	612
total sales	600	611	624	610	612
change in inventory	0	-11	-21	3	26
inventory	**-30**	**-41**	**-62**	**-59**	**-33**

Delay = 4, Gain = 0
Subject 10, Trial 1.1

In addition, she does not place enough emphasis on the role of change in inventory. In round 69, for example, she makes a production change of 0 despite the fact that round 68 revealed that there was a positive change in inventory and positive inventory.

round	68	69	70	71
change in production	12	**0**	-25	-80
production	668	668	688	706
dependent sales	0	0	0	0
independent sales	620	633	628	619
total sales	620	633	628	619
change in inventory	**48**	35	60	87
inventory	**15**	50	110	197

Delay = 4, Gain = 0
Subject 10, Trial 1.1

Sluggish control actions and ignoring the supply line result in large inventory discrepancies.

In game two (delay of 2 and gain of +0.3), subject 10 gives more attention to the change in inventory. For example, in round 72 she makes a zero change in production, noting from round 71 that change in inventory was negative. In addition, having realized that the supply line would wipe out the surplus inventory, she increases production in round 73 to compensate for oscillation.

round	68	69	70	71	72	73
change in production	-70	-60	-70	-80	**0**	**100**
production	578	528	458	398	328	248
dependent sales	137	119	98	74	74	104
independent sales	379	383	392	383	379	378
total sales	516	502	490	457	453	482
change in inventory	62	26	-32	-59	**-125**	-234
inventory	308	334	301	242	**117**	118

Delay = 2, Gain = +0.3
Subject 10, Trial 1.2

In game three (delay of 2 and gain of 0), subject 10 exhibits the same pattern of behavior as in the previous game emphasizing change in inventory. For instance, she begins game three with a surplus inventory and starts cutting production in fairly large increments.

round	62	63	64	65	66	67
change in production	**-30**	**-50**	**-50**	**50**	80	100
production	600	595	565	515	465	515
dependent sales	0	0	0	0	0	0
independent sales	569	562	561	561	574	577
total sales	569	562	561	561	574	577
change in inventory	31	33	**4**	-46	-109	-62
inventory	72	105	**109**	63	-46	**-108**

Delay = 2, Gain = 0
Subject 10, Trial 1.3

Her large decreases in production cause oscillations later (see inventory in rounds 66-67), forcing her to make large positive changes in production, thereby indicating weak supply line control.

Later games

In game one of session four (delay of 2 and gain of +0.6), we can see that she has not learned sufficiently the mechanics of delay nor gain. At first, she makes small decreases in production to decrease the growing surplus in inventory.

round	66	67	68	69	70	71	72
change in production	**-4**	**-4**	**-4**	**-4**	**-20**	**-20**	**-50**
production	596	596	592	588	584	580	560
dependent sales	355	353	350	348	336	324	294
independent sales	227	219	222	224	239	225	22[illegible]
total sales	582	572	572	572	575	549	518
change in inventory	14	24	20	16	9	31	42
inventory	**26**	**50**	**70**	**86**	**95**	**126**	**168**

Delay = 2, Gain = +0.6
Subject 10, Trial 4.1

When her decisions do not take effect right away, she increases the magnitude of changes and sets herself up with an oscillation situation, similar to other games, to which she must make large changes in production.

Effort

In terms of time spent on decisions, on average, subject 10 spends approximately 1.0 minutes per decision and shows several spikes of approximately 4 minutes in later rounds of each game. The later spikes did not correlate with strategy changes.

4.2.5 Subject 11

Early games

In game one (delay of 2 and gain of -0.3), subject 11 exhibits stage 3 strategy. He shows not only an appreciation of change in inventory but also a tacit understanding of the future change in inventory. For example, when both inventory and change in inventory are currently negative in round 62, he makes zero production change in round 63 instead of a positive change, indicating an appreciation of change in inventory. In the next round, although inventory is negative, he reacts with a decrease in production, showing that he understands that following rounds will increase in surplus.

round	60	61	62	63	64	65
change in production	0	**6**	20	**0**	**-3**	-10
production	600	600	600	606	626	626
dependent sales	-180	-182	-18[illegible]	-188	-187	-184
independent sales	783	796	791	780	791	804
total sales	603	614	603	592	604	620
change in inventory	-3	-14	**-3**	**14**	22	6
inventory	-3	-17	-20	**-7**	**15**	**21**

Delay = 2, Gain = -0.3
Subject 11, Trial 1.1

Subject 11 remains attentive towards the supply line throughout the game. Despite an inventory surplus of 139 in round 89, he increases production by 40 controlling for the insufficient level of production in the supply line.

round	88	89	90	91
change in production	-30	-30	**40**	20
production	599	569	539	509
dependent sales	-162	-153	-165	-171
independent sales	743	752	753	754
total sales	581	599	588	583
change in inventory	18	**-30**	**-49**	**-74**
inventory	169	139	90	15

Delay = 2, Gain = -0.3
Subject 11, Trial 1.1

In game two (delay of 0 and gain of 0), subject 11 focuses on inventory and change in inventory, and he exhibits a cautiousness bias. Because of the zero delay, attention to future change in inventory cannot be evaluated. It seems from his earlier game, with delay of 2, he has carried over his bias to remain cautious in his production control; therefore, he makes production changes with an "act and wait" style of heuristic. For example, in round 75 he waits until inventory is +40 before he initiates a significant control action to decrease production.

round	72	73	74	75	76	77
change in production	**0**	**0**	**-3**	**0**	**-20**	-20
production	565	565	562	562	542	522
dependent sales	0	0	0	0	0	0
independent sales	559	558	549	548	537	529
total sales	559	558	549	548	537	529
change in inventory	6	7	13	14	5	-7
inventory	6	13	26	**40**	45	38

Delay = 0, Gain = 0
Subject 11, Trial 1.2

In game three (delay of 0 and gain of -0.3), subject 11 shows a similar pattern to the one used in game two: sluggish control actions maintaining higher inventories. For instance, in round 68 to 73 he allows inventory to swell to +43, which in the 0 delay condition is unnecessary. The pattern of waiting for production changes to take effect possibly suggests that the subject learned in the first game to be cautious with production changes, and inappropriately transfers that strategy. This behavior might suggest that subjects develop cookbook rules and hold only fuzzy understanding for what situations to apply them.

round	68	69	70	71	72	73	74
change in production	**0**	**0**	**-5**	**-2**	**-18**	**-10**	**-20**
production	585	585	580	578	560	550	530
dependent sales	-176	-176	-174	-173	-168	-165	-159
independent sales	747	753	753	738	723	709	712
total sales	572	578	579	564	555	544	553
change in inventory	14	8	1	13	5	6	-23
inventory	**10**	**18**	**19**	**32**	**37**	**43**	20

Delay = 0, Gain = -0.3
Subject 11, Trial 1.3

Later games

In game three of session four (delay of 2 and gain of +0.6), his previous strategy which included a non-aggressive control of inventory combines with his delay control to cause inventory to remain in surplus throughout the game. For example, the game begins with inventory backlogged and swells to -13.0. He increases production to compensate, but inventory increases positively in the next round to +17.0, mostly due to a decrease in independent sales. Here, he begins to decrease production in response to a rapidly increasing inventory surplus. He continues to decrease production until he notices that change in inventory becomes smaller and then slightly negative. See below.

round	63	64	65	66	67	68	69
change in production	**-15**	**-20**	**-20**	**0**	**-30**	**-10**	0
production	600	615	600	580	560	560	530
dependent sales	360	348	336	336	318	312	312
independent sales	223	211	223	210	221	236	229
total sales	583	559	559	546	539	548	541
change in inventory	17	56	41	34	21	**12**	-11
inventory	**4**	**60**	**101**	**135**	**156**	**168**	**157**

Delay = 2, Gain = +0.6
Subject 11, Trial 4.3

The next rounds show a decrease in inventory, but he decides that the decrease is sufficient in magnitude and rate, as evidenced by five 0 production decisions in a row. Although inventory does decrease to 54, later it steadily increases in surplus until the end of the game. He never controls inventory aggressively enough.

Effort

Subject 11 makes almost no calculations throughout, while he does keep written track of independent sales.

Subject 11's times on decisions average approximately .5 minutes per choice. His decision times throughout were fairly consistent, indicating that he did not take much time to re-think strategies.

4.2.6 Subject 12

Early games

Game one (delay of 0 and gain of +0.3) reveals that subject 12 is very reflective of her strategy. She begins with one round of decreasing production, even though inventory begins backlogged. She then corrects this error, culminating her orientation phase, and makes small increases in production to decrease a growing backlog. It takes until round 72, about mid-game, before her increases in production bring backlog to near zero. In round 73, her decision of -4 shows some appreciation of the previous round's change in inventory that was positive 20 and larger than inventory. For example, see below:

round	69	70	71	72	73	74
change in production	**2**	**5**	**3**	**4**	**-4**	**-10**
production	631	636	639	643	639	629
dependent sales	189	191	192	193	192	189
independent sales	453	448	439	430	418	431
total sales	642	639	631	623	610	620
change in inventory	-11	-3	8	**20**	29	9
inventory	**-27**	**-30**	**-22**	**-2**	28	**37**

Delay = 0, Gain = +0.3
Subject 12, Trial 1.1

In round 69 she states, "I'm hanging even, but I've got to kill the backlog," indicating that she understands that her production increases need to be more effective at reducing the backlog. At round 71 she states, "Gradual changes seem the best since independent is so random that big changes cost a lot and have a high probability of over- or under-shoot." At round 74, when inventory has begun to show a surplus of +37, she says of her changes in production, " trying an experiment with high changes." At round 75 she says, "not good enough."

In round 89, despite the small magnitude of the production change, subject 12 shows further emphasis on change in inventory. She says, " I don't want to overshoot," indicating some cautiousness.

round	87	88	89	90
change in production	-17	-17	**2**	**6**
production	543	526	528	534
dependent sales	162	158	158	160
independent sales	371	383	392	393
total sales	534	541	550	553
change in inventory	9	**-15**	-22	-19
inventory	32	18	-5	-24

Delay = 0, Gain = +0.3
Subject 12, Trial 1.1

By the end of game one, this subject shows that she has learned to place greater emphasis on change in inventory. She reveals an ability to be reflective about her heuristics, but she concludes with a bias for smaller changes. She is one of the few subjects who indicates a willingness to explore alternative control actions.

In game two (delay of 2 and gain of -0.3), she continues the previous conservative changes in production and exhibits greater emphasis on change in inventory and future change in inventory. In rounds 80 and 81, for instance, she increases production in response to a negative change in inventory and previous decreases in production, i.e. future change in inventory, thereby exhibiting strategy 3. See below.

round	77	78	79	80	81
change in production	-8	-10	-20	**20**	**11**
production	556	549	541	531	**511**
dependent sales	-162	-159	-153	-159	-162
independent sales	709	696	710	701	699
total sales	547	537	557	542	536
change in inventory	9	12	**-16**	**-11**	**-25**
inventory	39	51	36	25	0

Delay = 2, Gain = -0.3
Subject 12, Trial 1.2

In game three (delay of 4 and gain of -0.3), subject 12 maintains the previous strategy 3. She begins by making more aggressive changes in production to compensate for a surplus inventory.

round	64	65	66
change in production	-7	-7	**8**
production	600	570	565
dependent sales	-166	-164	-166
independent sales	741	741	754
total sales	575	577	588
change in inventory	25	**-7**	-23
inventory	103	**96**	73

Delay = 4, Gain = -0.3
Subject 12, Trial 1.3

Therefore, she is showing greater attention to the inflection points of inventory and change in inventory. In later rounds, she decreases production in order to compensate for earlier increases, thereby showing an appreciation for delay; see below.

round	67	68	69	70	71	72	73
change in production	12	13	15	**-10**	**-3**	**-9**	**-30**
production	560	553	546	554	566	579	594
dependent sales	-170	-174	-178	-175	-174	-171	-163
independent sales	757	747	753	753	738	723	709
total sales	587	573	575	578	564	551	546
change in inventory	-27	-20	-29	-24	2	28	48
inventory	46	26	-3	**-27**	**-24**	**3**	**51**

Delay = 4, Gain = -0.3
Subject 12, Trial 1.3

Later games

In game 2 of session 4 (delay of 4 and gain of +0.3), subject 12 demonstrates a high level of sophistication and an appreciation for the effects of gain. Guessing that the game will end at round 92, subject 12 makes an increase in production of 400 to decrease surplus inventory, relying on the negative side effects of gain to bring inventory down.

round	88	89	90	91
change in production	0	**400**	0	0
production	463	413	413	413
dependent sales	124	244	244	244
independent sales	295	284	297	299
total sales	419	528	541	543
change in inventory	**44**	**-115**	**-128**	**-130**
inventory	516	401	274	144

Delay = 4, Gain = +0.3
Subject 12, Trial 4.2

The subject's ability to handle positive gain and delay have limitations, however. Just prior to the previous decision, she shows a reluctance to control inventory aggressively, indicating that her strategy is not adequate to deal appropriately with delayed, positive gain.

round	82	83	84	85	86	87
change in production	**-50**	0	0	**-50**	0	0
production	500	500	500	513	463	463
dependent sales	139	139	139	124	124	124
independent sales	346	333	320	305	297	293
total sales	485	472	459	428	421	417
change in inventory	15	28	41	84	42	46
inventory	231	259	300	384	426	472

Delay = 4, Gain = +0.3
Subject 12, Trial 4.2

Effort

In terms of time spent on decisions, on average, subject 12 had times of approximately .5 minutes. Her games generally showed lower times in the beginning rounds and a spike of approximately 1 -2 minutes in later rounds at roughly the middle of the game, which did not correlate with changes in strategy.

Game 4 of session 2 shows an increase in time over the other games of approximately 1 minute per choice. There were large spikes at rounds 80, of 10 minutes, and 82 and 83, of approximately 6 minutes, which did not correspond to changes in strategy.

4.2.7 Subject 13

Early Games

In game 1 (delay of 0 and gain of -0.6), subject 13 shows a focus on inventory and ignores change in inventory.

He begins by producing a surplus of 98 in response to an initial backlog of -35. There is no delay in this game, but the subject sets up large oscillations for which he must later respond. For instance, in round 68 he has an inventory of +13 and change in inventory of -268, which was produced by a change in production of -100 previously made. See below.

round	66	67	68	69	70
change in production	-5	**-100**	**-100**	-10	10
production	645	545	445	435	445
dependent sales	-387	-327	-267	-261	-267
independent sales	970	972	980	993	988
total sales	583	645	763	732	721
change in inventory	62	-100	**-268**	-297	-276
inventory	381	281	**13**	-284	-560

Delay = 0, Gain = -0.6
Subject 13, Trial 1.1

In rounds 83 and 84 of game one subject 13 makes some interesting changes in production. Rather than doing experiments with dramatic changes in production, i.e. whole numbers or hundreds, the subject makes changes in the *hundredths* and only expects consequences in the hundredths. Upon determining the effects of the hundredths change in production, the subject makes another extremely small change to further increase his knowledge of the system while incurring small cost, but the system, although it calculates at the thousandths precision, only shows the results in hundredths. The subject seems to design a rather creative no-risk experiment. It is reminiscent of small signal analysis and is a way to see effects without disturbing the system.

round	81	82	83	84
change in production	-20	-20	**-.01**	**-.0001**
production	180	160	159.99	159.99
dependent sales	-108	-96	-95.99	-95.99
independent sales	952	943	945	937
total sales	844	847	84[illegible].01	841.01
change in inventory	-664	-687	-689	-681.02
inventory	1098	411	**-278.02**	**-959.04**

Delay = 0, Gain = -0.6
Subject 13, Trial 1.1

Throughout game one, subject 13 remains in the orientation phase and does not show attention to change in inventory.

In game two (delay of 4 and gain of +0.3), subject 13's orientation phase continues for the initial rounds. For round 62 the subject writes parallel calculations:

```
Inv. 10 want to reduce inv.
prod = 600 x .3 => sales

590 * .3 = 177 (dep) + 410 = 587 total
change = 13
599.90 - 585.97 = 14

59.5 * 3 = 178.5 (dep) + 410.0 = 588.5 total
change      595.0 - 588.5 = 6.5
```

After two rounds of this type of experimentation, he abandons it for an inventory focus, making only changes that appear to be related to the magnitudes of inventory and change in inventory; shown below.

round	61	62	63	64	65	66
change in production	-0.1	-0.9	**-10**	**-40**	**-60**	**-100**
production	600	600	600	600	599.9	599
dependent sales	180	180	177	165	147	117
independent sales	406	410	402	397	389	388
total sales	586	590	579	562	536	505
change in inventory	14	10	21	38	64	94
inventory	24	34	**56**	**94**	**158**	**252**

Delay = 4, Gain = +0.3
Subject 13, Trial 1.2

Thus, through the above calculations, the subject reveals a focus on the difference between production and sales as being important. His subsequent control actions attempt to reduce the inventory gap to zero; however, he is still struggling with how to get inventory and change in inventory combined, at the same time remaining far away from understanding the supply line.

In game 3 (delay of 0 and gain of +0.3), he begins with the same experimentation with small production changes, but he immediately abandons this strategy and makes control changes in response to surplus inventory. By the end of the game, he is making smaller changes in production, and he is focusing more on change in inventory. For example:

round	86	87	88	89
change in production	-5	5	-1	5
production	523	528	527	532
dependent sales	156	158	158	160
independent sales	370	364	374	364
total sales	527	522	532	524
change in inventory	-4	6	-5	8
inventory	-7	-2	-7	2

Delay = 0, Gain = +0.3
Subject 13, Trial 1.3

Later games

The beginning of game 2 of session 2 (delay of 2 and gain of 0) reveals that subject 13 continues with his previous naive strategy of small production changes and has not learned from previous games how to control the system but only to be cautious. He starts the game with the same experimentation with extremely small production changes of $\pm$ .5 units. By round 64, the subject has stopped making his previous "what if" calculations and notes, for the first time, " change in inventory = production - sales total." In round 67 he states, "sales independent unpredictable, cut cost by keeping production constant," but he draws the wrong conclusion from this observation, as indicated by his later decisions that are overly cautious changes in production to large surplus inventory. The early behavior is depicted below.

round	62	63	64	65	66	67
change in production	**.5**	**-.5**	**0**	**0**	**.5**	**-.5**
production	600	600	600.5	600	600	600
dependent sales	0	0	0	0	0	0
independent sales	598	606	593	603	590	604
total sales	598	606	593	603	590	604
change in inventory	2	-6	7.5	-3	10	-4
inventory	3	-3	4.5	1.5	11.5	7.5

Delay = 2, Gain = 0
Subject 13, Trial 2.2

Later rounds of the same game show the subject's cautious heuristic. For example, in rounds 80 through 82, he makes very small changes in production despite large surpluses in inventory; he continues this behavior for the rest of the game.

round	80	81	82	83
change in production	**-5**	**-5**	**-7**	**-2**
production	570	565	560	555
dependent sales	0	0	0	0
independent sales	573	586	586	578
total sales	573	586	586	578
change in inventory	**-3**	**-21**	**-26**	**-23**
inventory	**226**	**205**	**179**	**156**

Delay = 2, Gain = 0
Subject 13, Trial 2.2

By game 3 of session 4 (delay of 4 and gain of +0.6), subject 13 remains overly cautious in his production changes and assumes a "no-control" posture by the end of the game. His control behavior in this game indicates that he believes that the best that he can do to control the system is to make very small changes in production to control cost. For example, in rounds 69 through 73, he makes production changes in -2 increments to a swelling surplus ranging from 135 to 235. See below:

round	69	70	71	72	73
change in production	**-2**	**-2**	**-2**	**-2**	**-2**
production	599	598	596	594	592
dependent sales	355	354	353	352	350
independent sales	229	224	211	217	220
total sales	584	578	564	569	570
change in inventory	15	20	32	25	22
inventory	**135**	**155**	**188**	**213**	**235**

Delay = 4, Gain = +0.6
Subject 13, Trial 4.3

By the end of the game, subject 13 has neither decreased inventory to zero nor risked having a backlog. By round 85, he makes a one time relatively large decrease in inventory, which does not decrease inventory significantly. He then makes no production changes and allows inventory to increase in surplus.

round	85	86	87	88	89	90	91
change in production	**-50**	**0**	**0**	**0**	**0**	**0**	**0**
production	507	497	487	477	427	427	427
dependent sales	256	256	256	256	256	256	256
independent sales	191	183	175	161	161	150	141
total sales	447	439	431	417	417	406	397
change in inventory	60	58	56	60	10	21	30
inventory	**770**	**828**	**884**	**944**	**953**	**974**	**1004**

Delay = 4, Gain = +0.6
Subject 13, Trial 4.3

Effort

Subject 13 makes a fair amount of "what if" calculations during the first five games. His early strategy, characterized by experiments with small production changes, does not prove to control the system sufficiently. In game 2 of session 2, he abandons his early strategy for cautious production changes to surpluses or backlogs. He continues the cautious, exogenous focus into the later games such that he assumes a no-control strategy.

With regard to time spent, subject 13 spends approximately .8 minutes per choice. His spikes are mixed between initial times of approximately 1.5 minutes and mid-game spikes of approximately 1.5 minutes. Thus, his times fit the pattern of having an initial strategy and then reevaluating it later.

4.3. Summary

Findings per subject

Subject 2 exhibits stage 3 strategy very early in the experiment and is to a large extent consistent throughout the experiment.

Although subject 4 attains a sophistication between strategy levels 2 and 3, he does not maintain a consistent strategy throughout the experiment. For instance, in several games, he develops a unique strategy for dealing with delay which results in undercontrol and requires him to switch strategies mid-game to strategies 1 and 2. His mathematical procedure miscalculates gain which results in large costs in the self-reinforcing conditions. Once the system is sufficiently perturbed by his mathematical procedure, he abandons his strategy and exhibits a mere inventory focus with control actions related to the magnitude of inventory: sophistication level 1.

Subject 8 exhibits strategy 3 by game three. His approach consists of minimal to no written calculations, and he is unique in keeping track of past values of production, inventory, and costs by graphing them. The extent to which he utilizes the graphs is not known; however, his strategy for dealing with the task is fairly good. Games 1 through 3 show a shift in emphasis on change in inventory and future change in inventory. Throughout there is a cautious approach, and the appreciation of delay is met with a cautious "act and wait" style.

Overall, subject 10 ignores delay in all of the games sampled and incorporates gain minimally. Subject 10 exhibits stage 2 strategy overall.

Subject 11 begins with and maintains a good understanding of how to control inventory under delay and gain conditions through monitoring change in inventory and future change in inventory closely, thereby exhibiting strategy 3 early on. His games show a cautious control bias.

Subject 12 exhibits stage 3 strategy and demonstrates incorporation of gain. However, she exhibits somewhat sluggish control changes to compensate for excesses or backlogs.

Subject 13 exhibits weak attention to change in inventory, but later games indicate a break down in his heuristics. Generally, he under-controls the system by ignoring gain and delay and making too small changes in production at the beginning of games which cause him problems and force him to make large changes in production later. In the final game, he remains cautious in his changes and assumes a "no-control" posture by the end of the game.

Group findings

Based on the above findings, the notebook analysis may be summarized as follows:

1. For almost all of the subjects, there is little orientation phase.

2. Most of the learning occurs within the first two games.

3. Most subjects seem not only to recognize the importance of inventory right away but also the importance of change in inventory.

4. Second games, for most subjects, reveal more emphasis placed on the change in inventory than initial games.

5. Increasing attention given to the change in inventory accounts for most of the performance improvement that occurs within the first three games.

6. Only some subjects show evidence of supply line control.

7. There seems to be a correlation between sophistication of strategy and scores achieved. Those subjects that attain stage 3 strategy, subjects 2, 8, 11, and 12, have better scores compared with subjects who attain only strategy levels 1 and 2, subjects 10 and 13. See figure 4.4

Subject	Log_2 Difference: Subject - Benchmark	Rank
12	-2.399	1
11	-2.124	2
8	-2.049	3
2	-2.025	4
3	-1.729	5
9	-1.534	6
4	-1.445	7
5	-1.003	8
7	-0.663	9
14	-0.483	10
10	-0.338	11
15	-0.048	12
6	+0.087	13
1	+0.909	14
13	+1.652	15

Figure 4.4 Subjects ranked according to $\log_2$ difference of costs from average of the optimum and no-control rules

In Figure 4.4 we provide subjects' ranked scores.[1]

There are differences in the attention given to supply line and gain, but it is difficult to assess the extent of the differences from this type of analysis. Subjects are able to adjust their heuristics only minimally according to varying conditions. All subjects exhibit a bias toward cautiousness as evidenced by 0 or small production changes and persistence of high inventory backlogs or surpluses under high delay and high gain situations.

The above findings help us to illuminate the three heuristic hypotheses proposed at the beginning of the thesis. Subjects appear to utilize more cues than the single-cue

[1] Score for subject = Log_2(cost of subject) - [Log_2(cost of optimum rule) + Log_2(cost of no-control rule)] /2 averaged over 15 trials (the geometric mean of the two benchmarks).

model would suggest. Rather subjects appear to select cues oriented toward expected change in inventory, which suggests a model more similar to the weight and cues adjustment heuristic. The rational model, on the other hand, does not appear to be supported from the findings, since subjects hardly ever consider all of the cues available, and they underweight the cues they do rely upon since their heuristics undercontrol the system under high delay and high gain.

Effort

Although not writing calculations does not necessarily mean not doing calculations, it seems that the number of calculations that the subjects write makes little difference in terms of scores. It is just as easy to perform many computations with little understanding as it is to perform the correct calculations. Most of the mental effort goes into the thoroughness with which subjects apply their strategy. For example, one of the subjects who writes the most in the notebooks, subject 13, performs many parallel "what if" computations and spends more time in performing these calculations than most subjects. However, he is also one of the worst subjects in terms of scores.

The sophistication and understanding of the system has almost nothing to do with the time spent; there are other factors that contribute to success. All of the effort spent in calculations will not help the subject if he or she has misperceived the fundamentals of the system (see subject 13).

Similarly, we may explain the practice effect described in the results section. The practice effect showed that subjects spent less time as the experiment progressed, even though their scores relative to optimal remained the same. The practice effect may be attributed to the subjects' understanding of the task, since after game three performance remains the same, and subjects merely get more efficient at executing their strategy.

The notebook results reveal that subjects' mental models do not contain sufficient consideration of high delay and high gain conditions. While not totally ignoring the supply line, it appears that subjects pay less than full attention to expected changes in inventory.

Informed by the previous analysis, the regression analysis will consist of building several increasingly sophisticated models, such as inventory only, inventory + change in inventory, inventory + expected change in inventory, and the rational (full state) model. Our notebook analysis would lead us to believe that most of the subjects fall into the two groups of inventory + change in inventory and inventory + expected change in inventory. The regression analysis should allow us to come to a more formalized understanding of the weight adjustment that subjects make to different information cues under different conditions.

5. Using adjusted R^2 to determine subjects' model

Introduction

The notebook analysis revealed that subjects went through a very brief orientation phase and that subjects' heuristics were fairly consistent across trials. While subjects' rules lead to satisfying results in the easier conditions, the same rules, under high delay and gain conditions, resulted in huge costs and eventually in an abandonment of the strategy employed. Without precise measurement of the changes in attention given to the various information cues in the task, it was difficult to assess from the notebooks the exact nature of the heuristics employed by the subjects. Thus, the notebook analysis provided qualitative insight into the cues and processing of subjects' heuristics. It was felt that the data required a more quantitative analysis, i.e. regression, that would measure cue weights and changes in weights over time, across conditions, and between different subjects. In addition, a regression analysis can provide a better indication of the extent to which high and low performers differ with respect to the emphasis given to particular cues.

5.1.1 Model specifications

In performing the linear regression analysis, four models of increasing sophistication were constructed which were informed both by the notebook analysis and the hypotheses that were stated at the outset of this thesis.

The models range from the least sophistication to the greatest sophistication.

$$\text{Model 1:} \quad \Delta\text{Production} = a1 * \text{Inventory}$$

Model 1 predicts that subjects' heuristics are based solely upon inventory values. Model 1 was not supported by the notebook analysis, but has been included to test the one cue adjustment hypothesis and as a benchmark for comparison with other rules.

$$\text{Model 2:} \quad \Delta\text{Production} = a1 * \text{Inventory} + a2 * (\text{Production-Sales})$$

Model 2 consists of two cues, inventory and change in inventory, implying that subjects look at both values of inventory and of the change in inventory, i.e. production - sales, in formulating their decisions. Subjects would use the change in inventory from the round previous to the current decision round combined with the value of inventory to determine an appropriate change in production. The notebook analysis revealed that some, although not the majority of the subjects might have followed this rule and it is used to test the weights only hypothesis.

$$\text{Model 3:} \quad \Delta\text{Production} = a1 * \text{Inventory} + a2 * (\text{Production[future]-Sales})$$

Model 3 is composed of inventory and future change in inventory, i.e. future production - sales, which includes whatever delay is present. Model 3 would predict that subjects make production changes that are based upon the values of inventory and consideration of the supply line through future changes in inventory. For instance, in a 2 delay situation, the subject would consider inventory and change in inventory two rounds into the future to arrive at the production decision.

Model 3, arising from the notebook analysis, would fall under the weight and cue adjustment heuristic, despite the fact that the number of cues remains the same. As delay increases, subjects' consideration of cues that reach into the future widens. While we did not anticipate this, we can subsume it under the weight and cue adjustment.

Full information model: ΔProduction = a1 * Inventory +
a2 * Sales +
a3 * Production +
a4 * Production[one round ahead] +
a5 * Production[two rounds ahead] +
...

Although the notebook analysis did not support the full information model, previous research has proposed this model as viable and, therefore, it was included in this regression analysis. The full information model is the most sophisticated heuristic which includes all of the production states in addition to the previous cues (inventory, sales, and future production-sales). This model assumes that subjects consider all of the possible information cues.

It is important to mention that the models were tested both with and without the constant term (a of the equation y = a + bx). In neither case were the fits significantly better, i.e. the t-values for the constant terms were not significant. This is expected since there should not be a bias towards a positive or negative inventory.

Adjusted R^2 and the number of sign reversals were used as indicators of the fit of a particular model. Several of the models tested resulted in signs of coefficients that would be counterintuitive. For instance, it makes no sense at all to increase production in response to a positive inventory. Thus, a positive weight on inventory would be judged an artifactual sign reversal. In no case did the notebook analysis even hint slightly that subjects would perform this way. Thus, we categorized a regression that showed one or more wrong signs on the cues considered as a case of sign reversal. Figure 5.1 shows the average adjusted R^2 and number of sign reversals. Information cues considered are presented in the rows, model specifications at the bottom of the columns. The average adjusted R^2 and number of sign reversals for the fifteen conditions can be found at the top of each column above the model name. These results will be discussed further in the following section.

Information cues considered	M1	M2	M3	M4
Average adjusted R^2	.26	.51	.58	.62
Number of sign reversals	24	59	24	136
Inventory	X	X	X	X
Production(t)-Sales		X		
Production(farthest)-Sales			X	
Production(t)				X
Production(t+1)				X
Production(t+2)				X
Production(t+3)				X
Production(t+4)				X
Sales(t)				X

Figure 5.1 Average adjusted R^2 and sign reversals for alternative models

5.1.2 Detailed sign reversal analysis

Model 1

Figures 5.2 - 5.5 show the number of correct signs per condition and the total number of sign reversals overall out of the possible 225 correct signs.

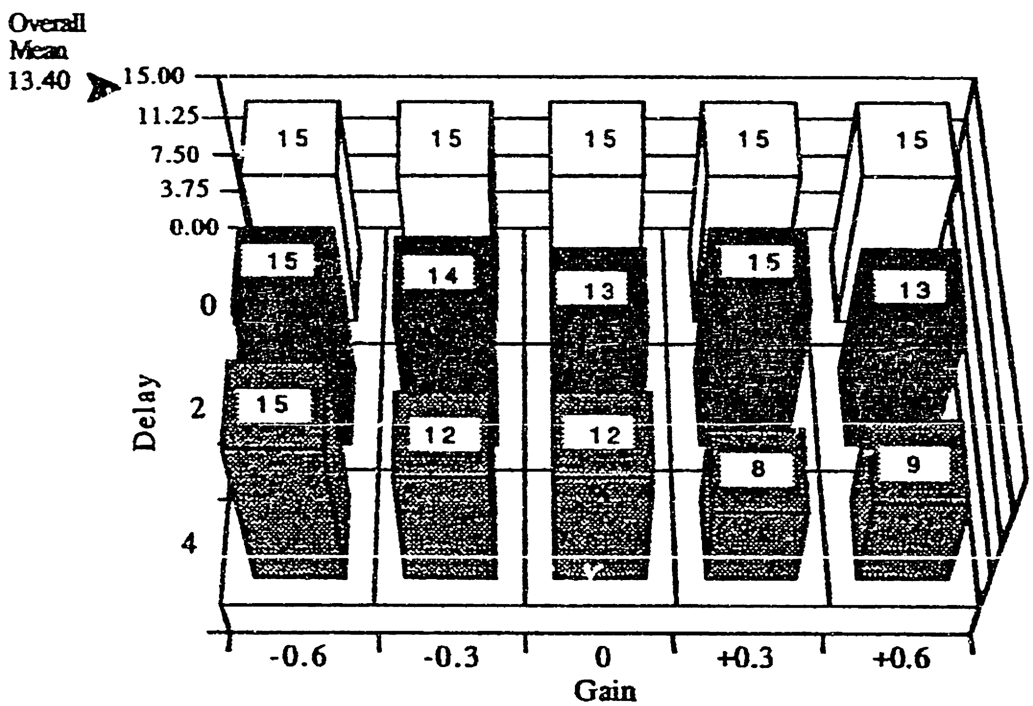

Total number of regressions with at least one sign reversal: 24 (11%)

Figure 5.2 Model 1: Number of instances where no coefficient of fit of the model has a wrong sign

Figure 5.2 shows that model 1 had 24 sign reversals of the total 225. Sign reversal frequency remains relatively constant across conditions, but decreases with increasing delay and gain.

Model 2

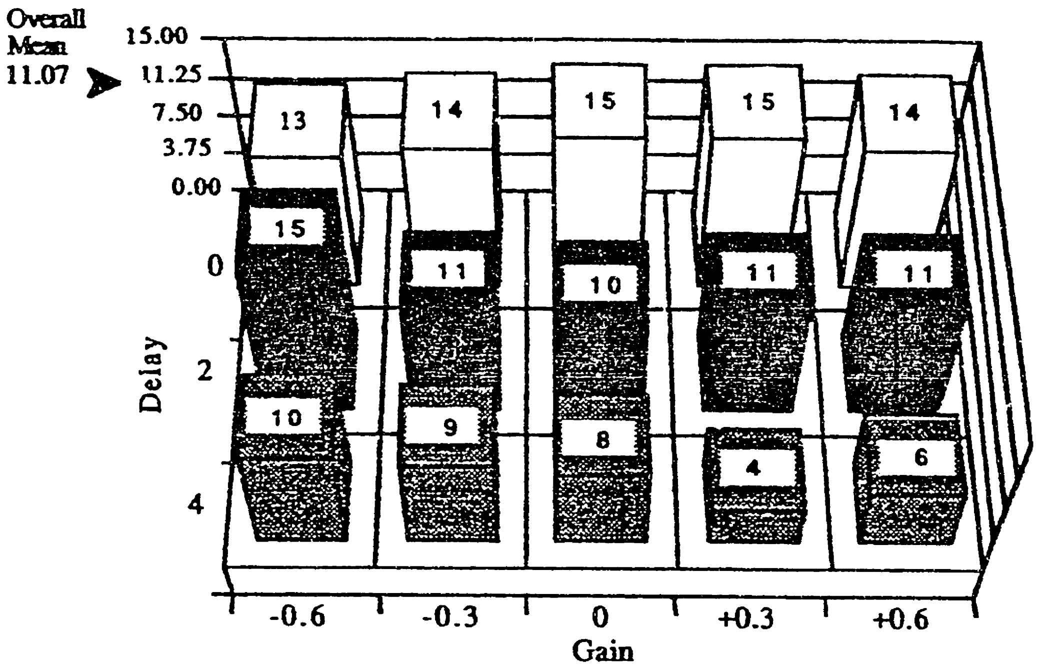

Total number of regressions with at least one sign reversal: 59 (26%)

Figure 5.3 Model 2: Number of instances where no coefficient of fit of the model has a wrong sign

Figure 5.3 shows that, compared with model 1, model 2 has increased sign reversals to 59. Sign reversals increase with increasing difficulty.

Model 3

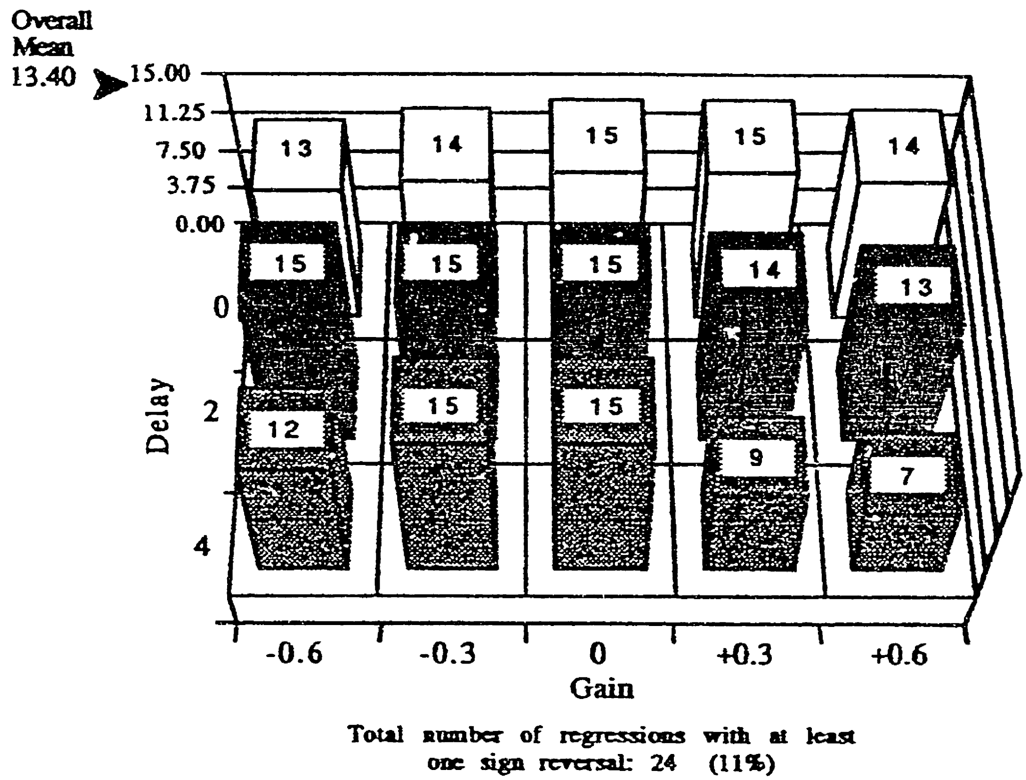

Figure 5.4 Model 3: Number of instances where no coefficient of fit of the model has a wrong sign

Figure 5.4 shows that model 3 had 24 sign reversals which is lower than model 2 and equal to model 1.

Full Information Model

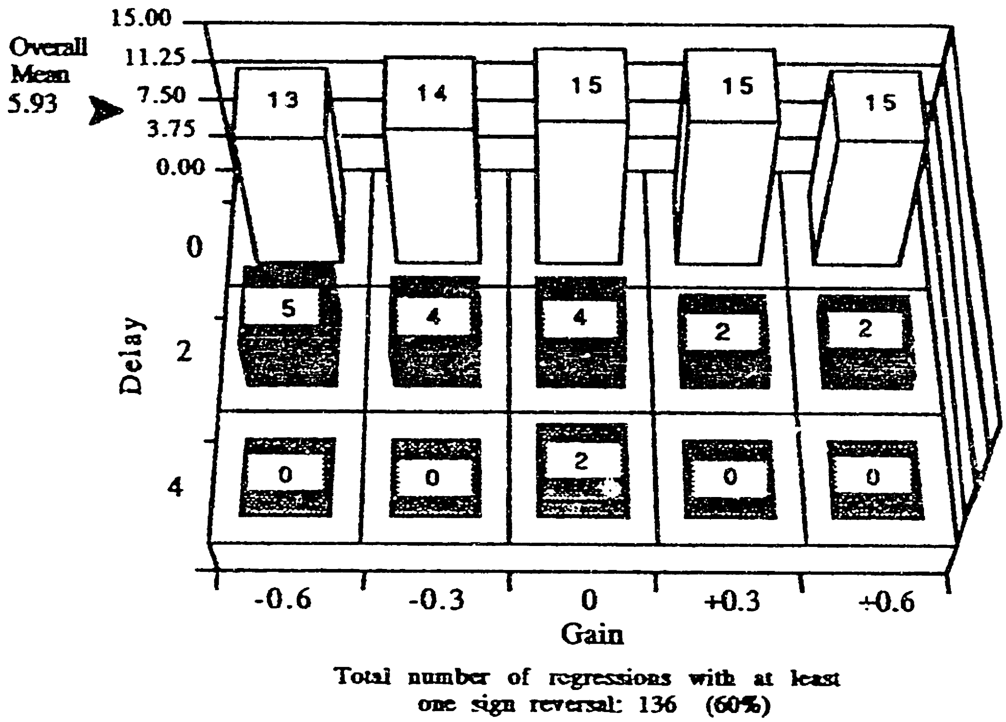

Figure 5.5 Full information model: Number of instances where no coefficient of fit of the model has a wrong sign

Figure 5.5 shows that the full information model had 136 sign reversals which is extraordinarily high compared with the other models.

Besides the four models discussed in this section, seven models ranging in sophistication between models 3 and 4 were also tested. These models consiste of adding various states intermediate between model 3 and the complete state model in order to test whether subjects' behavior might fall in sophistication somewhere between model 3 and the full-state model. All of the alternative models fell in the close range of R^2 of .58 and .62; also, all of these models produced sign reversals substantially higher than model 3. While we can clearly distinguish between models 1, 2, and 3, R^2 alone is unable to differentiate between models 3 and 4. We judged the number of sign reversals for the full information model extremely high, and we felt that the use of more sophisticated statistical techniques, such as restricting the signs on coefficients, would not help us with the basic dilemma of sharply distinguishing between models 3 and 4.

Kleinmuntz (1990, p.9) points out that at some point regression analysis is not a fine enough procedure to distinguish varying models of subjects' heuristics.

> "...the linear model can be viewed primarily as a description of the information *combination* process, and does not address the process of information *search* and *acquisition* (Einhorn, Kleinmuntz, & Kleinmuntz, 1979). Observing that the supply line coefficient (aSL) is too small does not tell us what the decision maker was actually thinking about. For instance, did the decision maker simply think the supply line was irrelevant, or was the problem that the decision maker had difficulty remembering or reconstructing what was in the supply line? Linear models do not provide any way to distinguish between these two possibilities.
>
> Thus, the linear model provides a very limited amount of insight into just what is going on in the decision maker's mind. One way to characterize the situation is that inferences from the linear model are useful in *identifying* the phenomenon of misperception of feedback, but they do not provide an entirely satisfactory *explanation* of why or how the phenomenon arises."

We find the regression analysis useful to clearly reject models 1 and 2; in the notebook analysis we found that the majority of the subjects learned to attend to future changes in inventory. However, regression becomes misty as more cues are considered, and we must rely primarily on the notebook analysis to judge model 3 as the most appropriate model to describe subjects' behavior.

5.1.3 Detailed R^2 Analysis

We not only examined the overall means but, suspecting that different subjects might have used different models and that there may have been variable fit across trials, we also analyzed subject and trial differences. Thus, the standard ANOVA toolkit used earlier was applied.

Previous research suggests that performance in decision tasks depends on the adequacy of the decision maker's mental model. High-performers and low-performers might not only differ in the weight they assign to different information cues, but also in the selection of those cues. If the latter is true, we would expect to see differences in model fit across subjects. Once we have identified the most likely model used by the subjects, the next section will explore the weights applied to cues.

To test the adjusted R^2 for treatment, practice and subject effects, an ANOVA was performed. The results for each of the 4 models are documented in 5 tables using the standard toolkit as introduced in previous sections: (1) Average adjusted R^2 across conditions; (2) Adjusted R^2 across conditions showing variability; (3) Adjusted R^2 across trials; (4) Adjusted R^2 across subjects in predetermined order; and (5) the results of the ANOVA.

Model 1

Overall performance

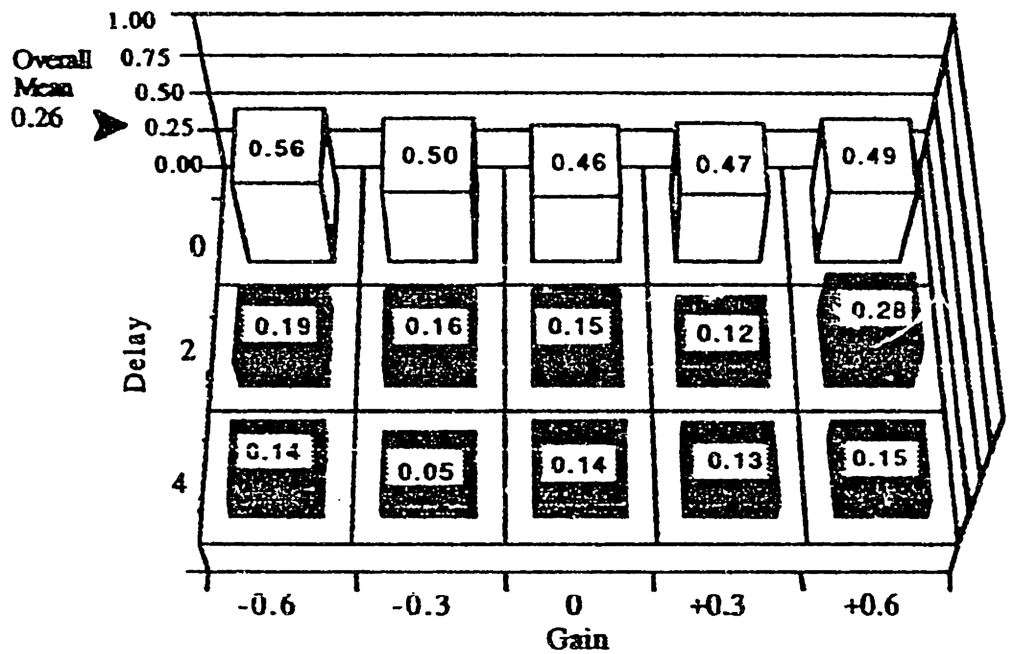

Figure 5.6 Model 1: Average adjusted R^2 across treatments

Model 1, which consists of only one cue, inventory, produced an average R^2 of .26. See Figure 5.6 which shows the average R^2 across conditions and overall. Under zero delay, model 1 provides a fit of approximately .50. The fact that fit decreases as delay increases can be taken as an indication that people are considering more cues than inventory alone.

Analysis

Source	df	F	probability
Overall treatment	14,182	17.40	$p<.01$
Delay	2,182	114.05	$p<.01$
Gain	4,182	1.91	------
Delay x Gain	8,182	0.99	------
Subject	14,182	2.86	$p<.01$
Practice	14,182	1.26	------

-- not significant

Figure 5.7 Model 1: ANOVA summary table

ANOVA Figure 5.7 provides the statistical analysis of R^2 and reveals that the differences in fit across treatments is attributable to delay (which reflects the number of relevant cues).

Treatment effects

Figure 5.8 shows R^2 across treatments with variances. Variance remains fairly high across conditions, but decreases slightly as gain and delay increase. The cause for the high variance is due to differences between subjects who themselves differed in performance across conditions.

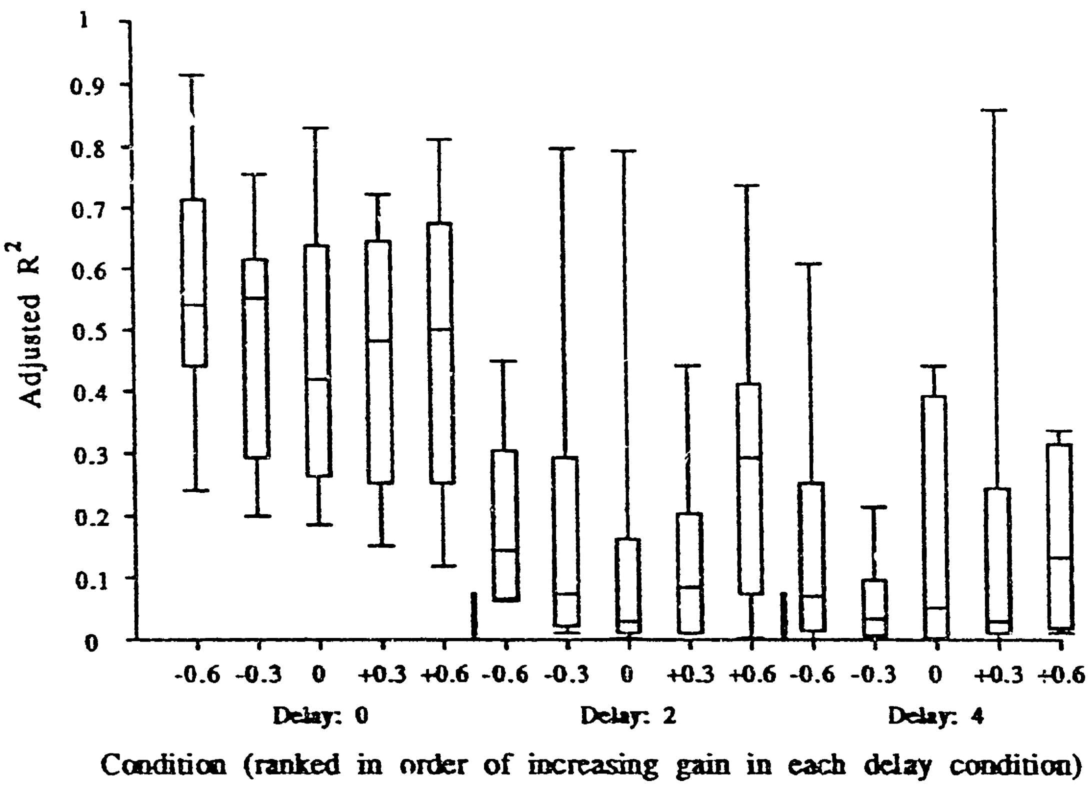

Figure 5.8 Model 1: Adjusted R^2 variance across treatments

Practice effects

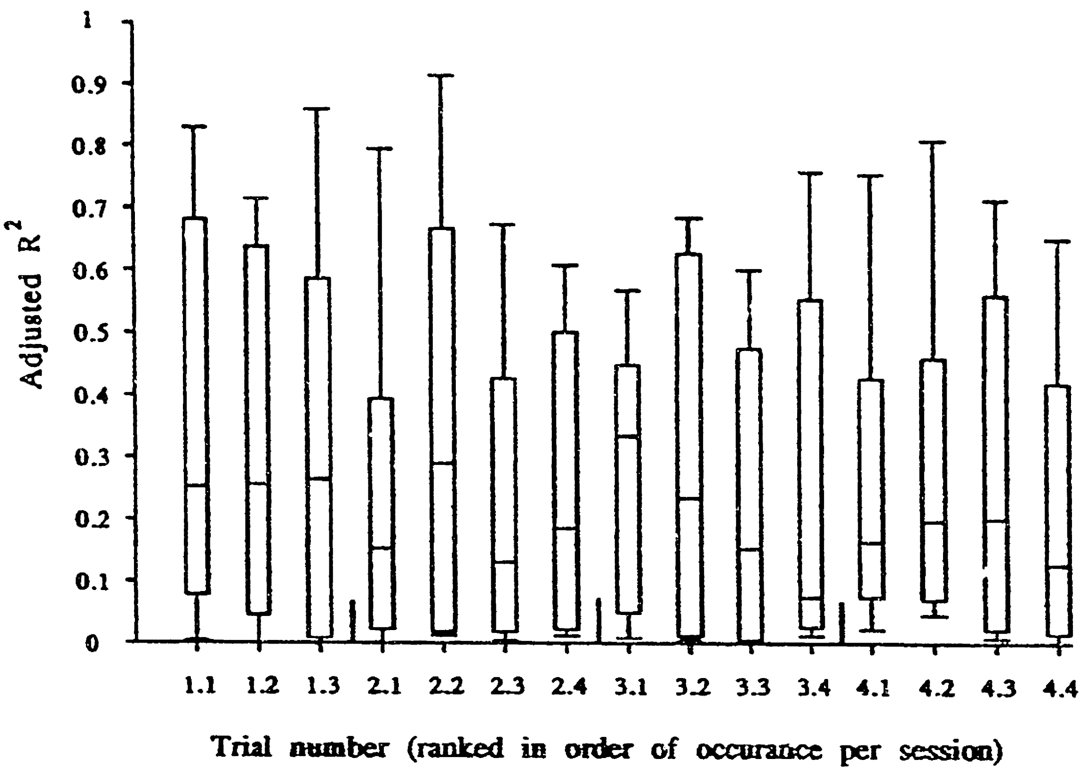

Figure 5.9 Model 1: Adjusted R^2 variance across trials

From ANOVA Figure 5.7, practice effects are not significant and can be seen in Figure 5.9.

Subject effects

R^2 for subjects did vary significantly and can be seen in Figure 5.10. Model 1 explains subjects 6, 10, 7, and 13 best with an average R^2 of approximately .30. This result is expected for subject 10, since the notebook analysis revealed that she attained a shaky sophistication level 2, attention to inventory and change in inventory because her attention to change in inventory was weaker than many other subjects. Subject 13's heuristics followed a similar pattern as subject 10, thus model 1 is consistent with these subjects' abilities.

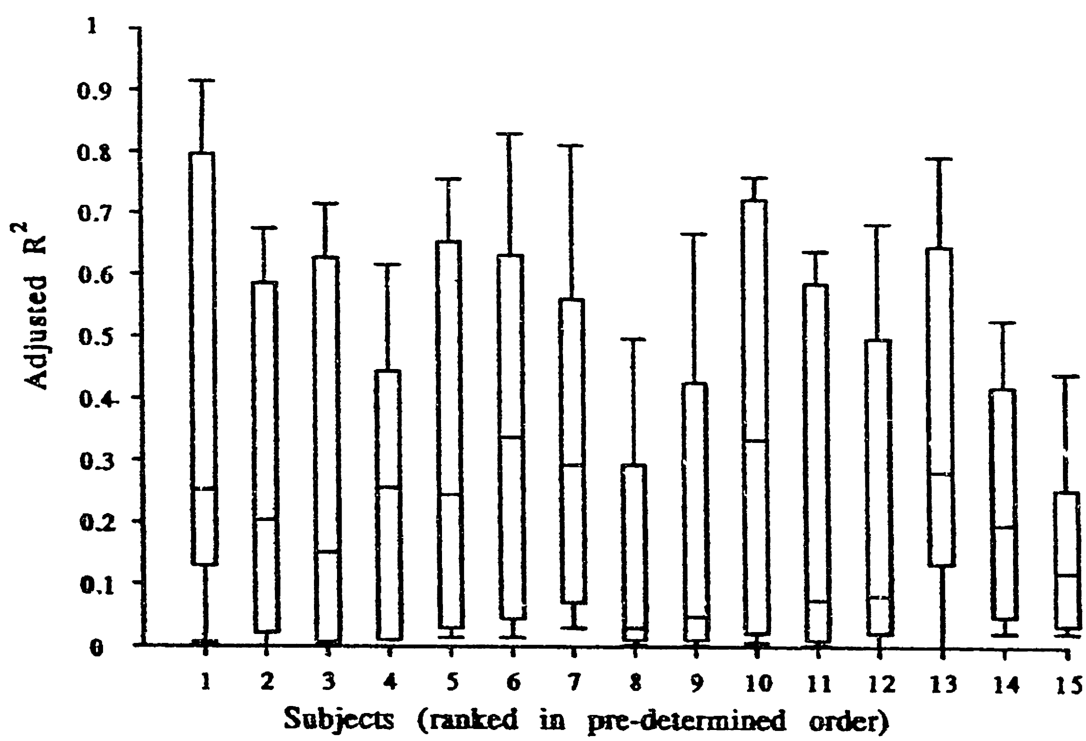

Figure 5.10 Model 1: Adjusted R^2 variance across subjects

Model 2

Overall performance

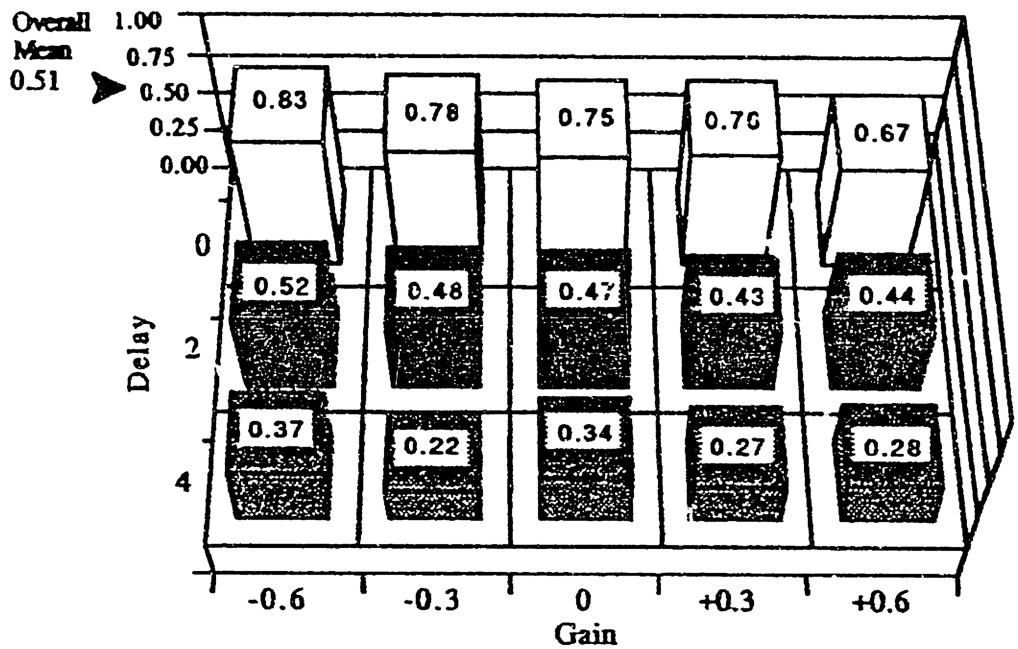

Figure 5.11 Model 2: Average adjusted R^2 across treatments

Model 2 assumes that subjects consider change in inventory (production -sales) in addition to inventory discrepancy. As Figure 5.11 shows, this model accounts for .51 of the variance of subjects' decisions. Under zero delay, model 2 accounts very well for subjects' decisions with an average R^2 of approximately .75. With increasing delay, the relative fit decreases as hard conditions produce more variance. Nevertheless model 2 accounts for more variance than model 1, as indicated by higher R^2.

Analysis

Source	df	F	probability
Overall treatment	14,182	15.21	p<.01
Delay	2,182	100.56	p<.01
Gain	4,182	1.87	------
Delay x Gain	8,182	0.55	------
Subject	14,182	1.75	p<.05
Practice	14,182	0.26	------

-- not significant

Figure 5.12 Model 2: ANOVA summary table

ANOVA Figure 5.12 indicates that there were differences in fit across treatments. The fit across treatments is attributable to delay.

Treatment effects

Figure 5.13 reveals the above in more detail and shows that across treatments, the variance remains roughly the same. Variability is explained as differences between subjects.

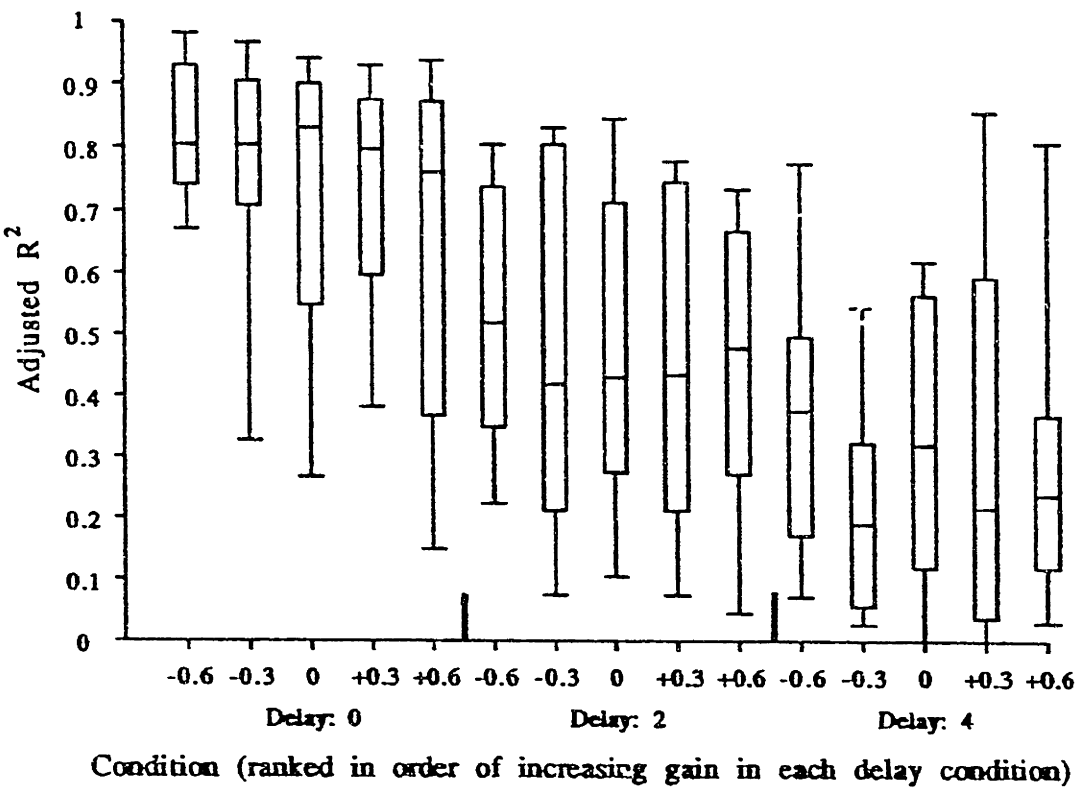

Figure 5.13 Model 2: Adjusted R^2 variance across treatments

Practice effects

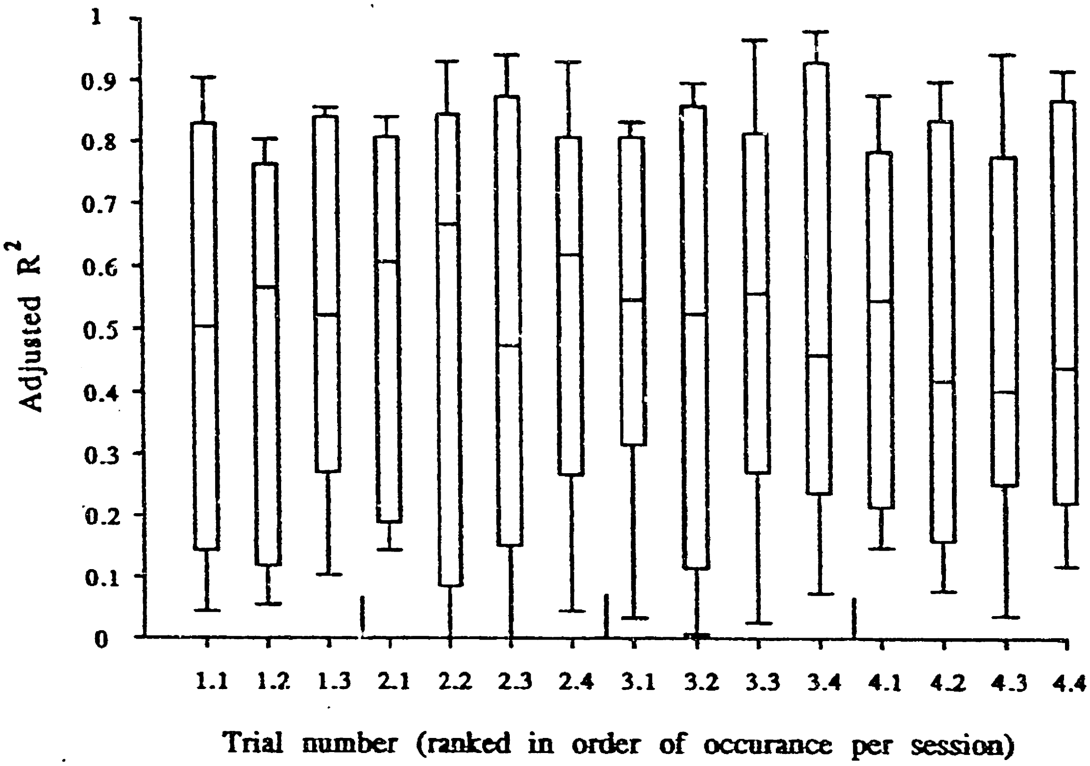

Figure 5.14 Model 2: Adjusted R2 variance across trials

ANOVA Figure 5.12 shows no effect due to sequence of trials, see Figure 5.14.

Subject effects

Subjects differ with respect to fit by model 2. As Figure 5.15 shows, overall, subjects' decisions are accounted for with an average of approximately .40. Subjects 2, 7, and 10 show the highest fit for model 2 with an average of approximately .60.

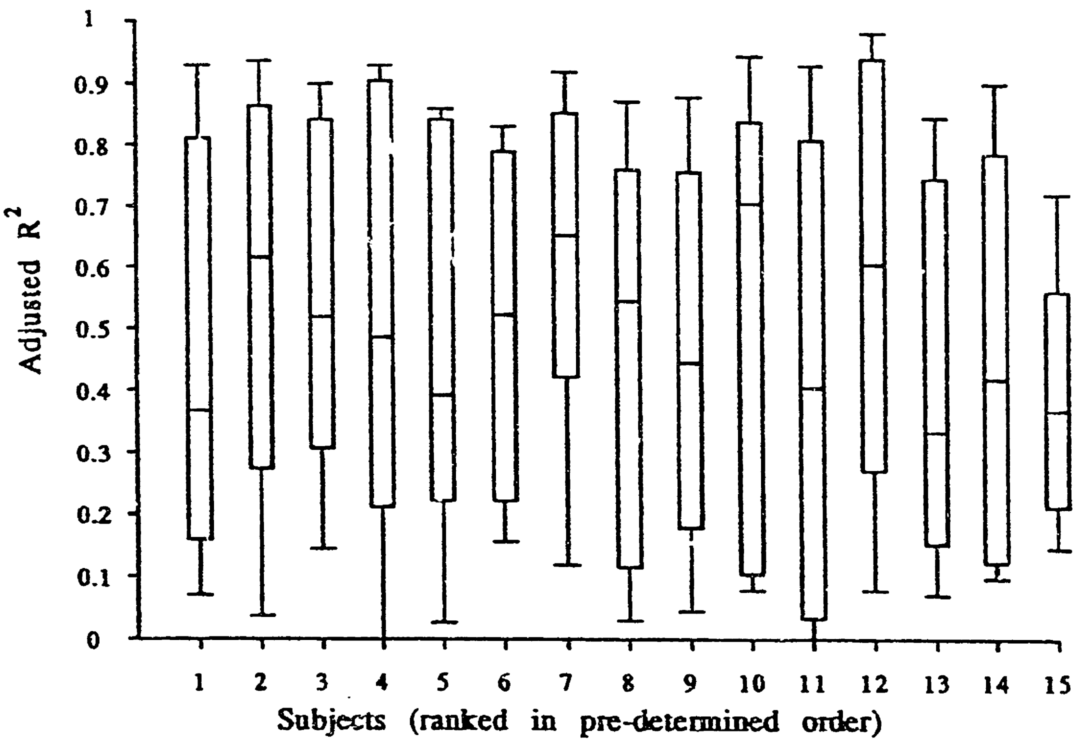

Figure 5.15 Model 2: Adjusted R^2 variance across subjects

The fact that this model accounts better for subjects' performance is not surprising, since the notebook analysis revealed that most subjects attained sophistication level 2. The fit to subject 10's data, in particular, is approximately .70, since she remained at a weak level 2 consistently throughout the experiment. The .3 fit for subject 13, on the other hand, is surprising, since he too attained only strategy level 2. We will see that increasing sophistication of models does not provide a better fit for subject 13, indicating that he remained at a sophistication level between inventory and change in inventory and was inconsistent.

Model 3

Like model 2, Model 3 consists of two cues. However, where model 2 considers current change in production, model 3 is more sophisticated in that it looks to future change in production as future productions appear on the horizon. Recall from the notebook analysis it appeared that most subjects not only considered current change in inventory but also anticipated future changes in inventory as it would change two and four weeks into the future.

Overall performance

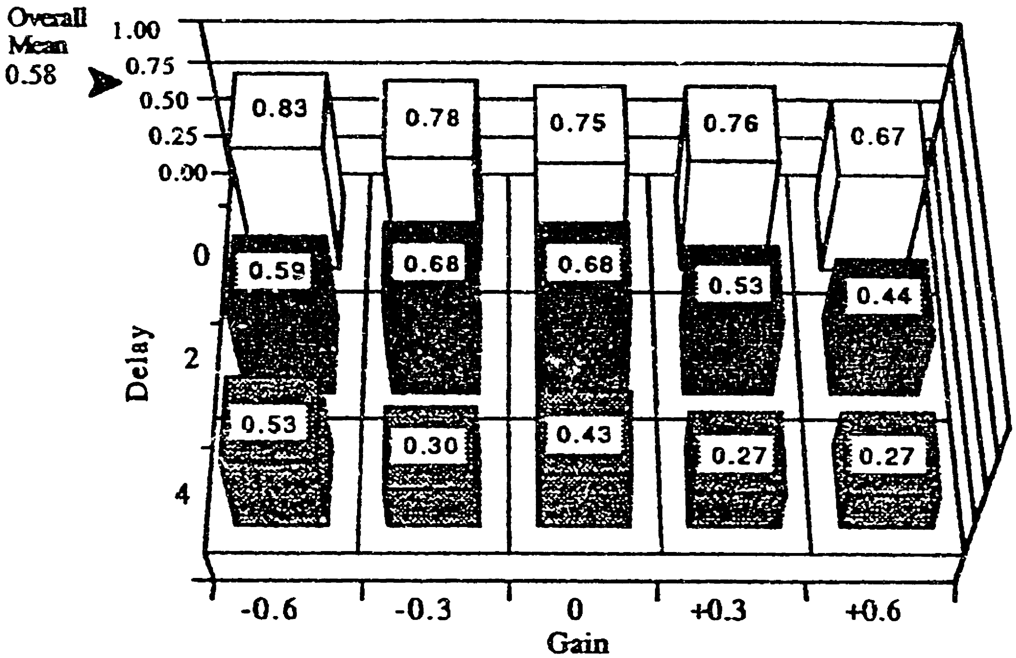

Figure 5.16 Model 3: Average adjusted R^2 across treatments

Figure 5.16 shows the average R^2 per condition and overall. The average is .58. Model 3 and 2 are identical in fit under the zero delay case, since there are no future production states to be considered. Both delay 2 and delay 4 cases are better accounted for with model 3 than model 2. In addition, recall from the sign reversal analysis that model 3 holds fewer sign reversals than model 2, e.g. 24 sign reversals versus 59, respectively. Thus, the overall judgement is that for most of the subjects model 3 is a better description than model 2.

Analysis

ANOVA Figure 5.17 shows that treatments are significantly different for the R^2 and mostly attributable to delay, although an effect of gain and an interaction is present. Subjects differ as well.

Source	df	F	probability
Overall treatment	14,182	16.36	$p<.01$
Delay	2,182	89.97	$p<.01$
Gain	4,182	8.22	$p<.01$
Delay x Gain	8,182	2.02	$p<.05$
Subject	14,182	2.66	$p<.01$
Practice	14,182	0.48	------

-- not significant

Figure 5.17 Model 3: ANOVA summary table

Treatment effects

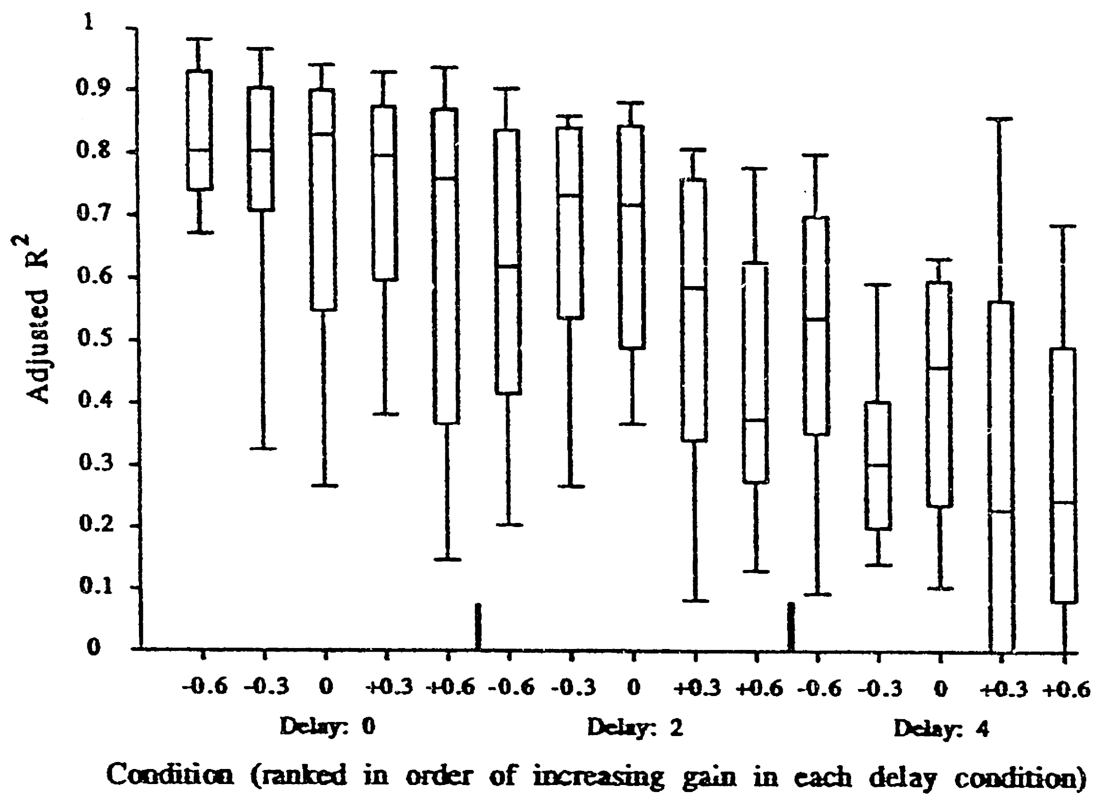

Figure 5.18 Model 3: Adjusted R^2 variance across treatments

Figure 5.18 provides a look at variance across conditions. Variance increases beyond the easiest condition, but does not increase much afterward. As described above, R^2 does decrease with increasing delay and gain. It is evident that subjects' heuristics break down under more difficult conditions; the reasons behind the decreased R^2 could be that subjects apply their rules with less consistency or they may simply abandon their initial strategies that worked well in easier conditions. The breakdown in heuristics is corroborated by the notebook analysis.

Practice effects

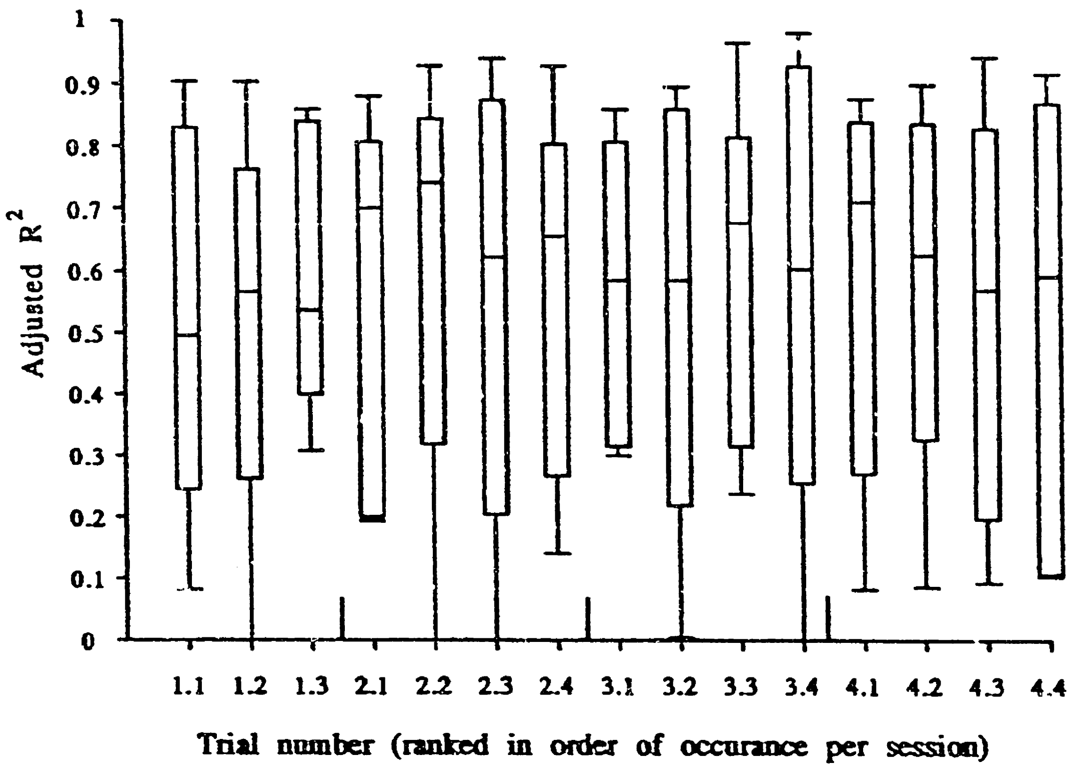

Figure 5.19 Model 3: Adjusted R^2 variance across trials

ANOVA Figure 5.17, again, shows no differences in fit due to the sequence of trials, revealed in Figure 5.17.

Subject effects

Subject differences are apparent. Figure 5.20 shows that model 3 accounts for most subjects' heuristics extremely well. Subject 1 and 13 are exceptions. Subject 1 was not examined in the notebook analysis; subject 13 showed a relatively low consistency of strategy across games. In fact, at the end of the later game of high delay and gain, he assumed the "no-control" posture. The lack of a consistent strategy and the generally low level of sophistication would explain a low R^2 for this particular subject. Furthermore, model 1, R^2 approximately .3, did not attain a better fit for subject 13 than model 3, R^2 of roughly .3 also.

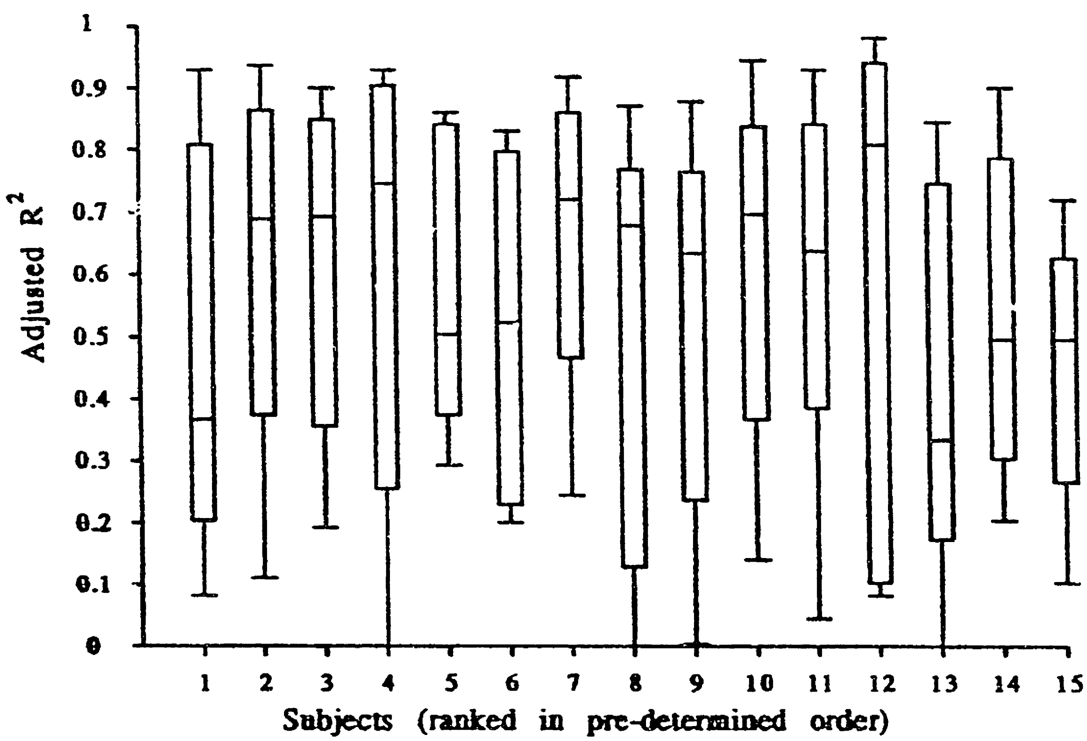

Figure 5.20 Model 3: Adjusted R^2 variance across subjects

Summary of analysis

We have seen a significant increase in R^2 in model 3. We feel comfortable in having demonstrated for the majority of the subjects that they achieved model 3 sophistication. Most of the subjects were sophisticated enough to take change in inventory into account and sophisticated enough to look into the future instead of just the present. These regression analyses corroborate nicely with the notebook analysis. The regression analysis does not support a further sophistication level in subjects, such as considering all of the states. The regression tool was not sharp enough. We were faced with a situation in which additional states placed in the model increased adjusted R^2 only minimally, but at the same time produced artificial sign reversal results.

Since we did not find any indications for higher sophistication from the notebook analysis, the rest of this analysis will use model 3. At the same time, we feel a little uneasy about how adequately model 3 fully captures the delay 4 case under most difficult conditions.

5.2. Weights attached to Cues

In the previous section we attempted to identify the cues that subjects consider in their decision making process. Having identified those cues as inventory and future change in inventory, we will use the current section to analyze what weights subjects associate with those cues. While we already know from section 3.3 that subjects systematically undercontrol the system, we do not yet know if they underweight all cues equally or if there are particular biases to be discovered. With this perspective in mind we will analyze Model 3 that the notebook analysis and section 4.2., the R^2 analysis, filtered out as the most likely decision model. Recall from the R^2 analysis that the constant term was determined to be not significant in describing subjects' performance. We will discuss model 3 with respect to both the coefficients that subjects attached to cues and the t-values that these cues attained in order to show how particular cues are emphasized and whether weights on those cues are significantly different from zero.

5.2.1 Inventory

Overall performance

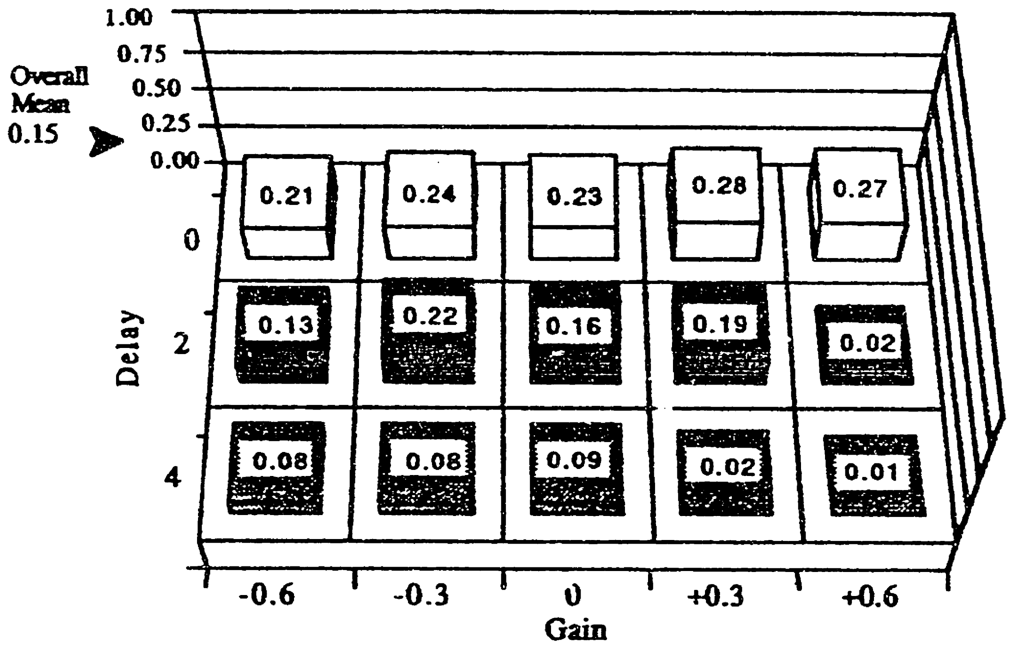

Figure 5.21 Average inventory weights: across treatments

Figure 5.21 shows the average weights assigned to inventory across treatment conditions. On average, subjects assign a weight of 0.15 to inventory. Across treatments, the weight assigned to inventory decreases as delay increases. Under zero delay, the average weight is approximately .25; under 2 delay, the average weight drops to approximately .14 and shows greater variability across gain conditions. Under 4 delay, the average weight drops lower to approximately .05 and shows a decrease across gain conditions. Subjects' performance will be contrasted with the weights assigned by the optimal rule in the following section.

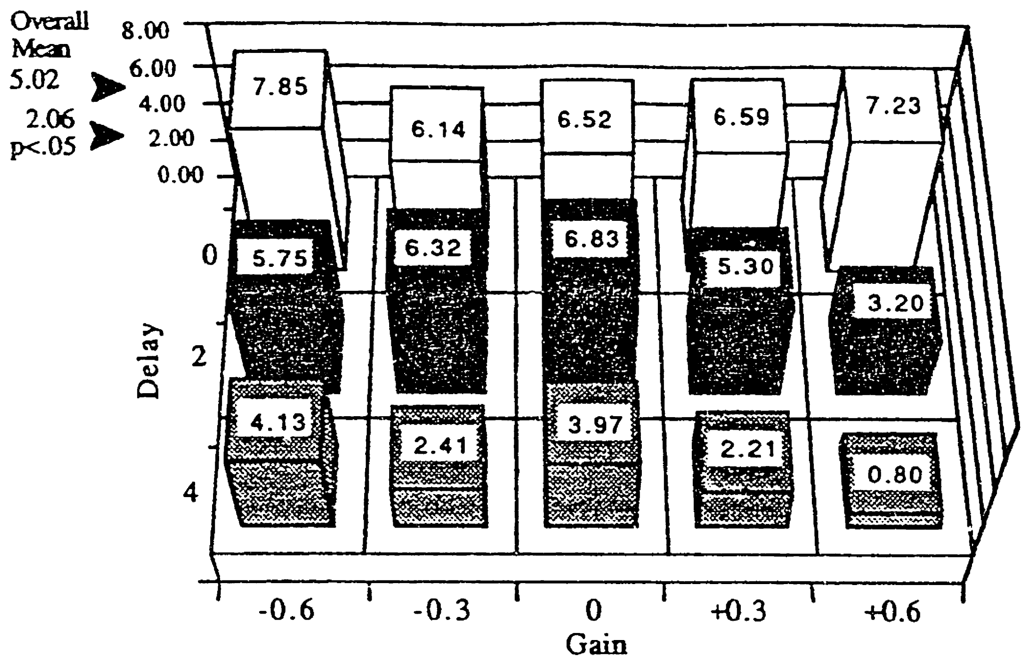

Figure 5.22 T-values for inventory weights: across treatments

The average t-values for inventory are presented in Figure 5.22. T-values above 2.06 are significant ($p<.05$). Fit decreases with difficulty. It seems that subjects might be operating with different rules in these cases. However, as we saw from the control effort section, subjects significantly undercontrol the system under increased delay and gain compared with the optimal rule. The decreased fit here for the high gain and delay corroborates that story and shows that subjects might abandon their own rule when faced with difficult circumstances.

Treatment effects

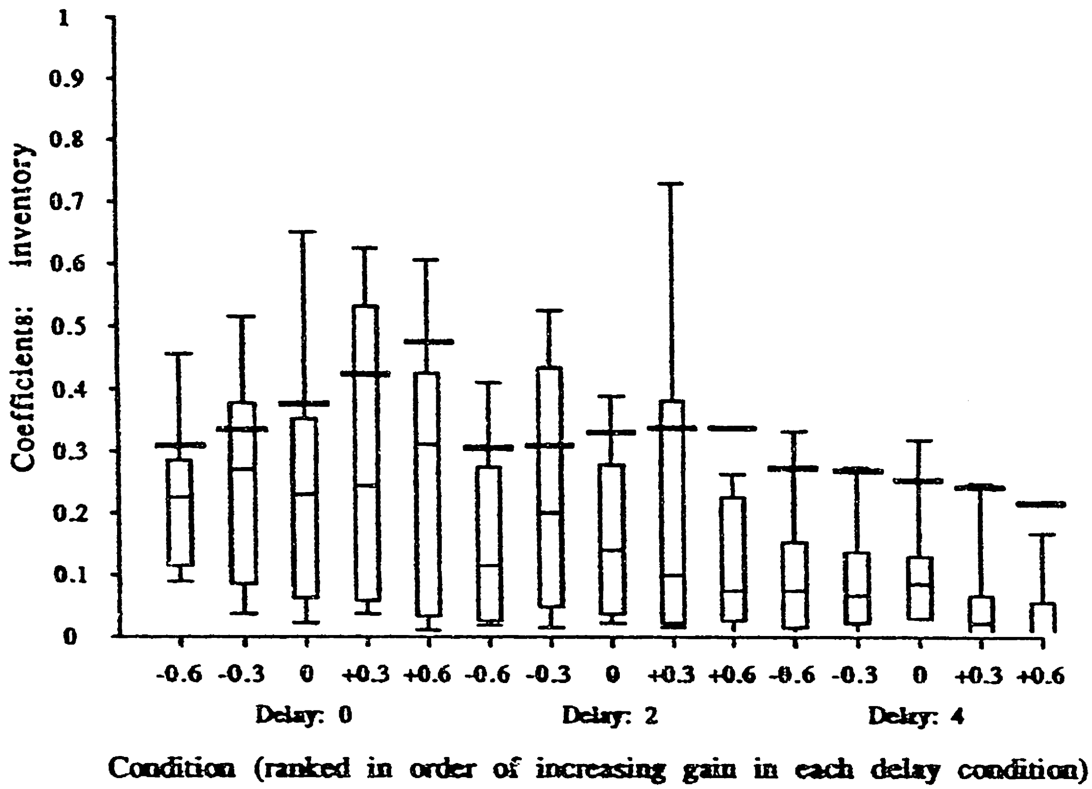

Figure 5.23 Variance of inventory weights: across treatments

In Figure 5.23, we see the same pattern from above: subjects' weights assigned to inventory decrease as delay increases. Optimum weights for inventory are shown as the thick, black bars. While the optimum rule associates roughly the same weights to inventory as delay increases, subjects' decrease that weight.

Notice that variability decreases with increasing delay. It seems that the weights assigned to inventory decrease in particular for those subjects who in easier conditions utilized relative high weights.

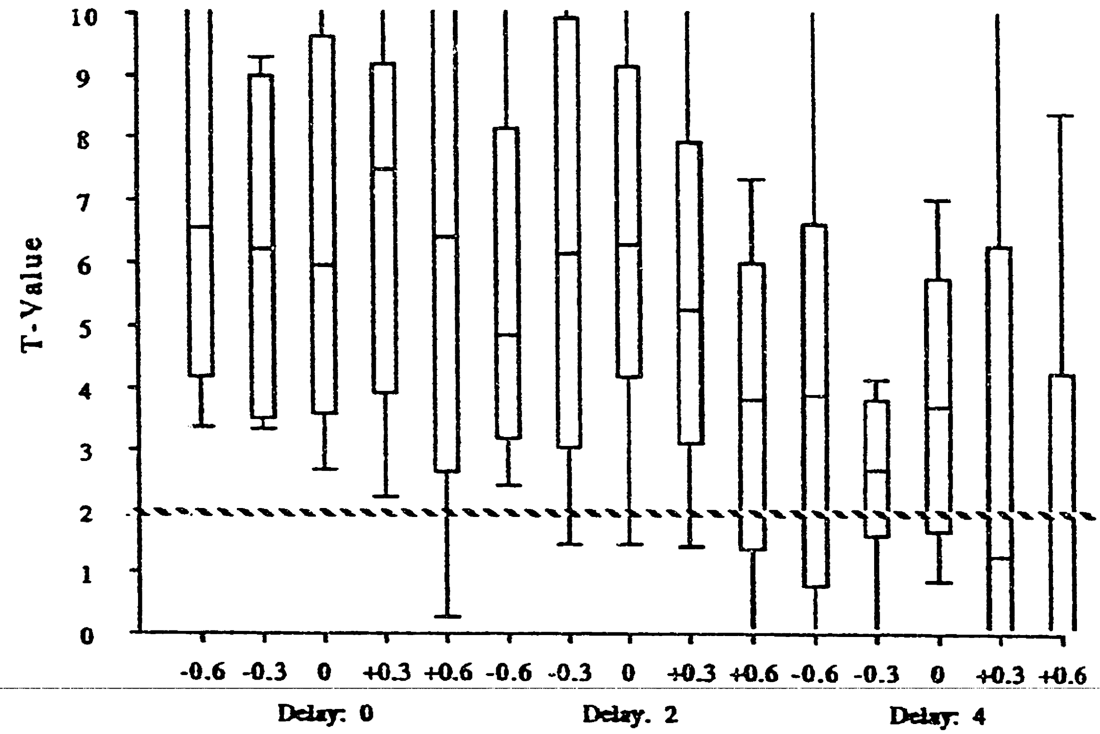

Figure 5.24 T-values for inventory weights: across treatments

Figure 5.24 shows the variance of t-values for the different conditions. T-values above 2.06 are significant, ($p<.05$). Fit decreases as delay increases, and variance of fit increases with increasing positive gain. In the 4 delay, positive gain cases, the estimated coefficients do not prove significant for the majority of the cases.

Practice effects

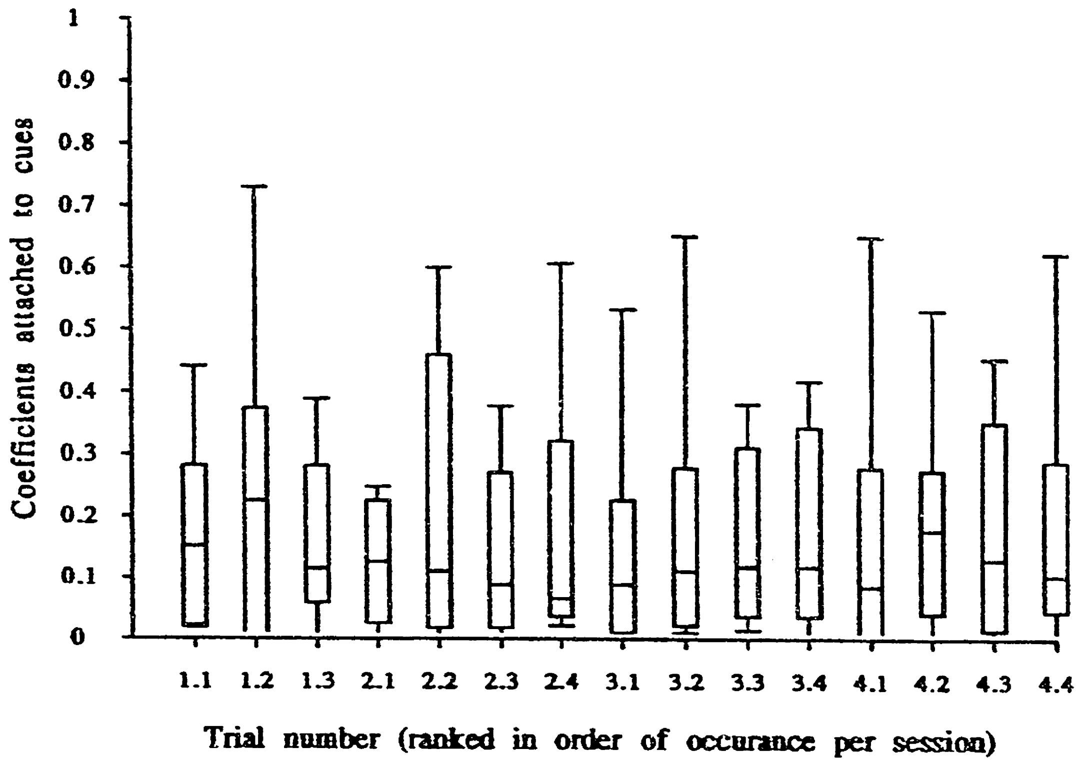

Figure 5.25 Variance of inventory weights: across trials

The sequence of trials does not have an effect on the weights assigned to inventory as Figure 5.25 and 5.26 shows.

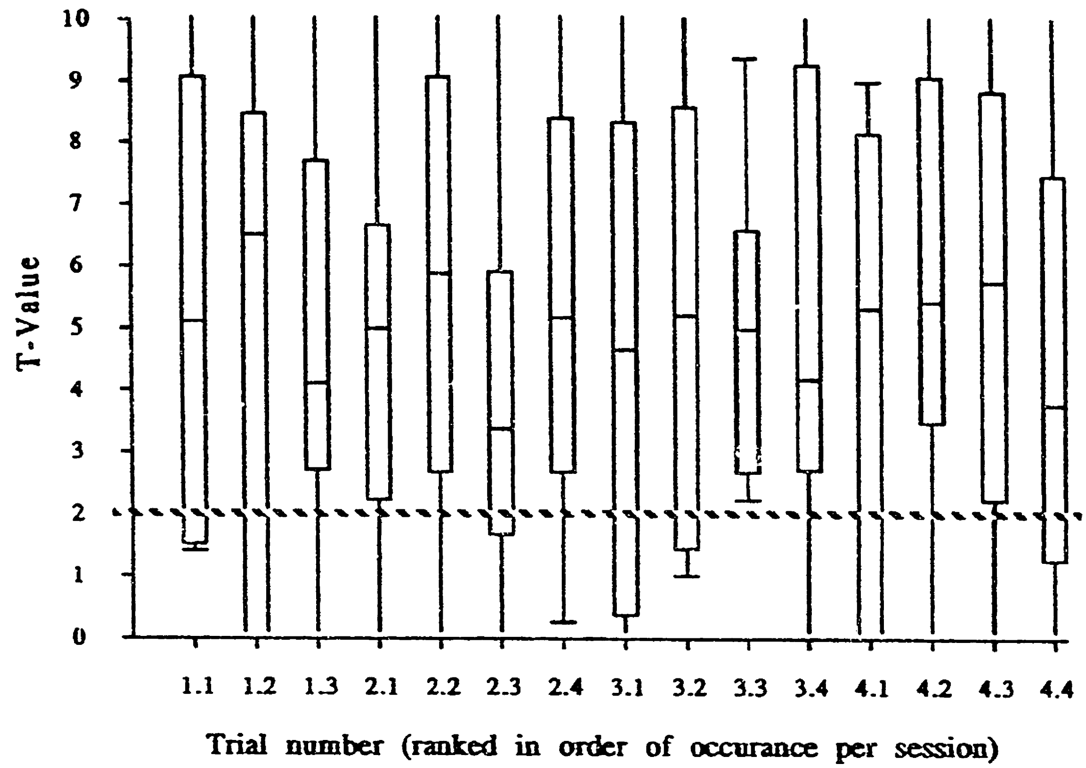

Figure 5.26 T-values for inventory weights: across trials

Subject effects

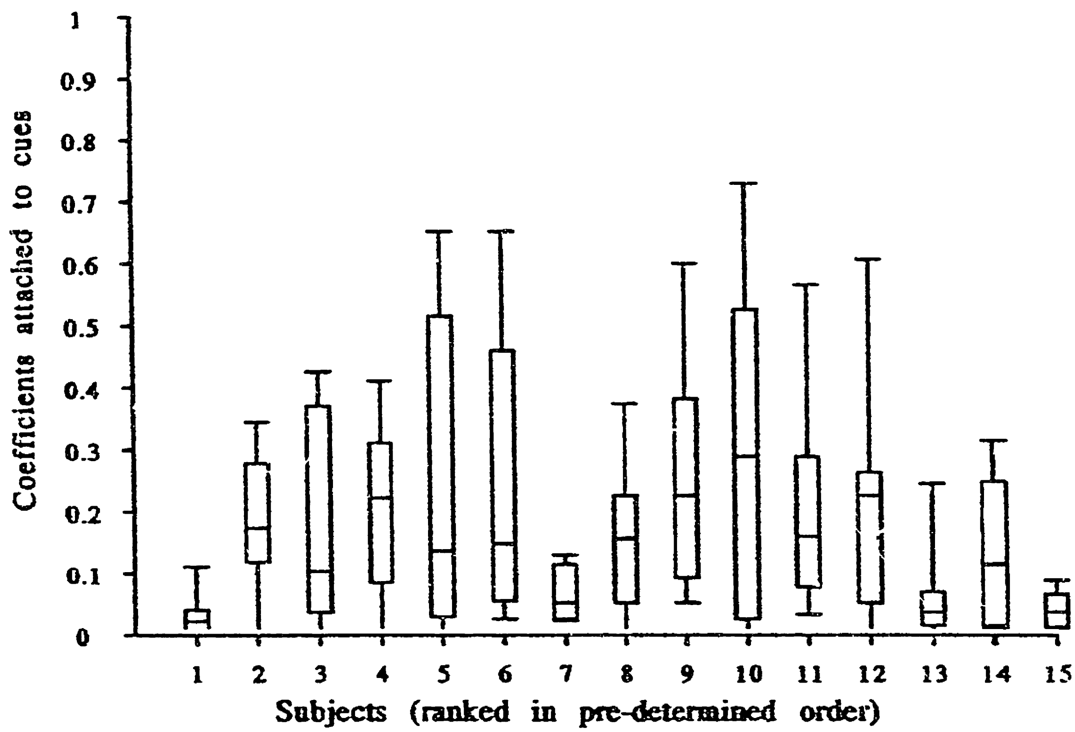

Figure 5.27 Variance of inventory weights: across subjects

Figure 5.27 reveals that different subjects weight inventory differently. In particular, subjects 1, 7, 13, and 15 place very little emphasis on inventory and do so consistently as indicated by the small variability, while the rest of the subjects weight inventory around .2 but do so with much variability.

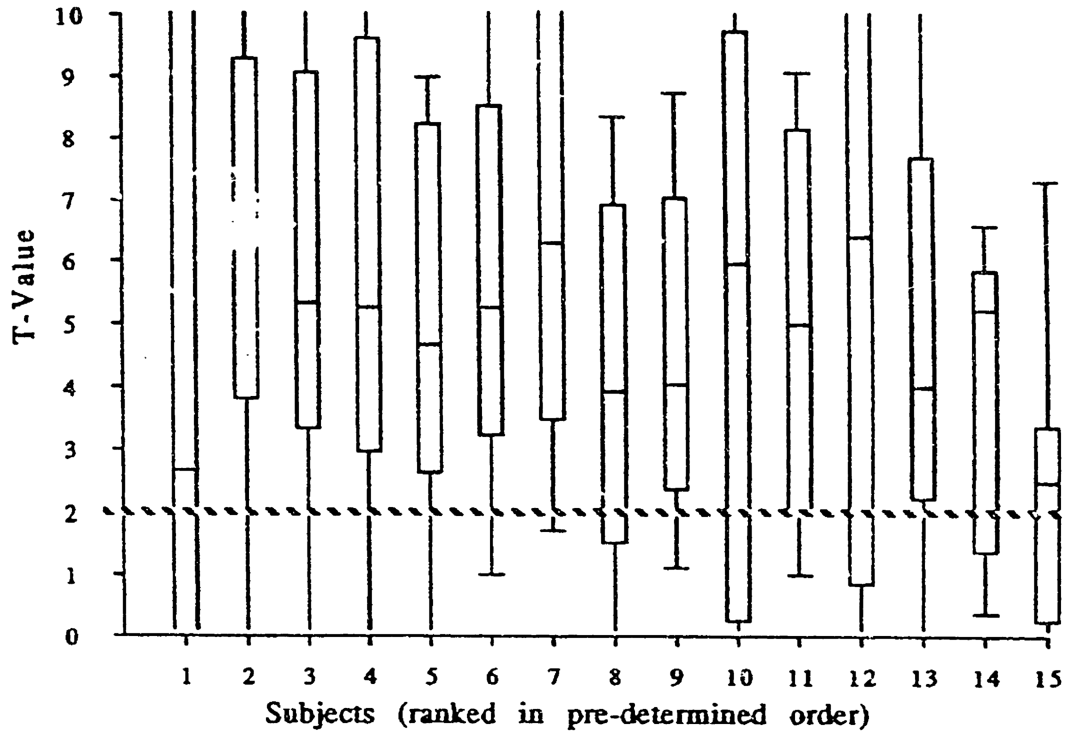

Figure 5.28 T-values for inventory weights: across subjects

Figure 5.28 shows the t-values for the subjects. Note that the majority of the t-values are significant for all subjects. This holds true also for those subjects who placed very low weight on inventory.

5.2.2 Future change in inventory

Overall performance

The weights associated with the future change in inventory stay roughly the same across conditions except for the 4 delay condition. The difference between the highest and lowest weight is approximately 250%, while the difference in the inventory case was much greater.

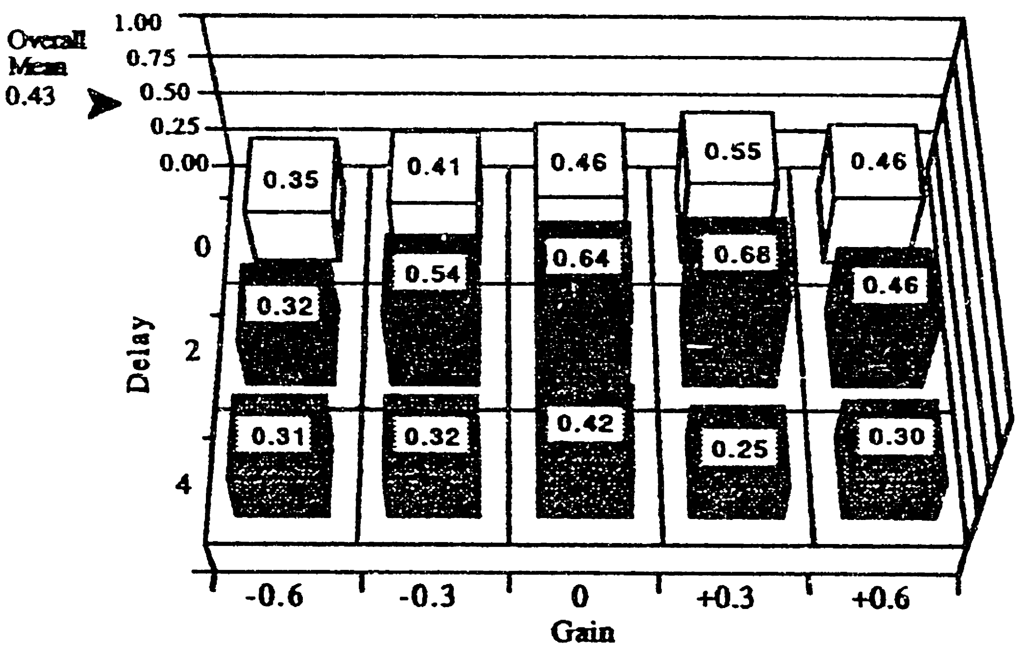

Figure 5.29 Average (future production - sales) weights: across conditions

Figure 5.29 shows the average coefficients applied to future change in inventory across treatments. The mean is .43. Figure 5.30 reveals that for most of the conditions, excluding the most difficult, the weights determined for future change in inventory are significant at least at the ($p<.05$) level. Fit decreases for increasing delay and gain, implying perhaps that subjects apply less weight to this cue as difficulty increases. From the notebook analysis, we saw that indeed subjects' heuristics broke down under increasing difficulty, and that effect can be described more precisely as a decrease in attention given to important cues.

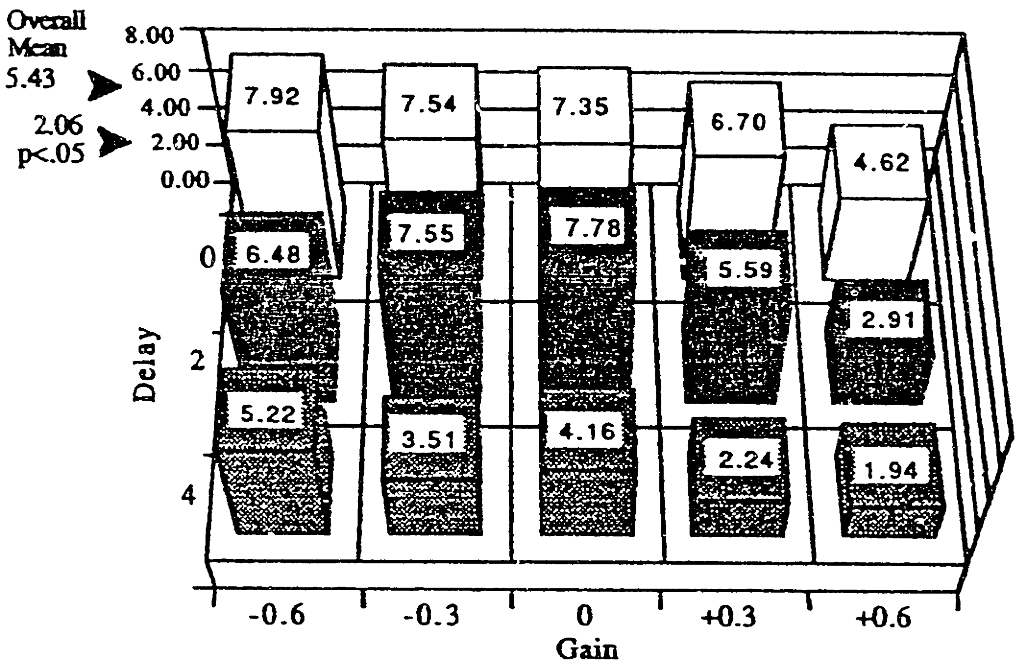

Figure 5.30 **T-values for (future production - sales) weights: across conditions**

Treatment effects

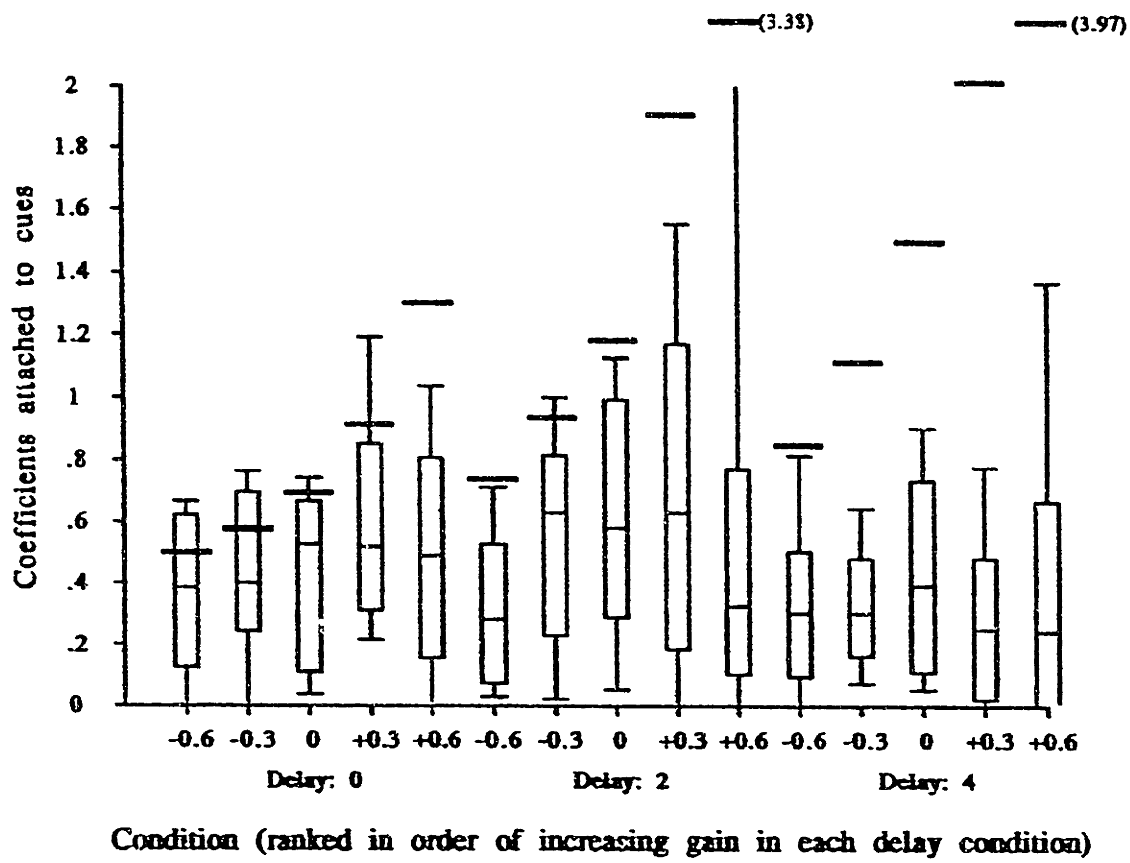

Figure 5.31 Variance of (future production - sales) weights: across conditions

Figure 5.31 shows the variance of coefficients assigned to future change in inventory. Optimum weights are shown as thick, black bars. While subjects stay relatively the same, optimum control suggests a very different behavior. Subjects' weights increase in the appropriate direction as long as complexity does not become too great. Once the difficulty; complexity of the task increases due to the combination of high delay and high gain, the task becomes too difficult and subjects' weights fall toward zero, possibly indicating a "hands off" control posture. In the difficult conditions, subjects' weights approach 0 instead of being more aggressive. These results support the notion that subjects' weights for this cue decrease with increasing difficulty, and especially for high delay and gain.

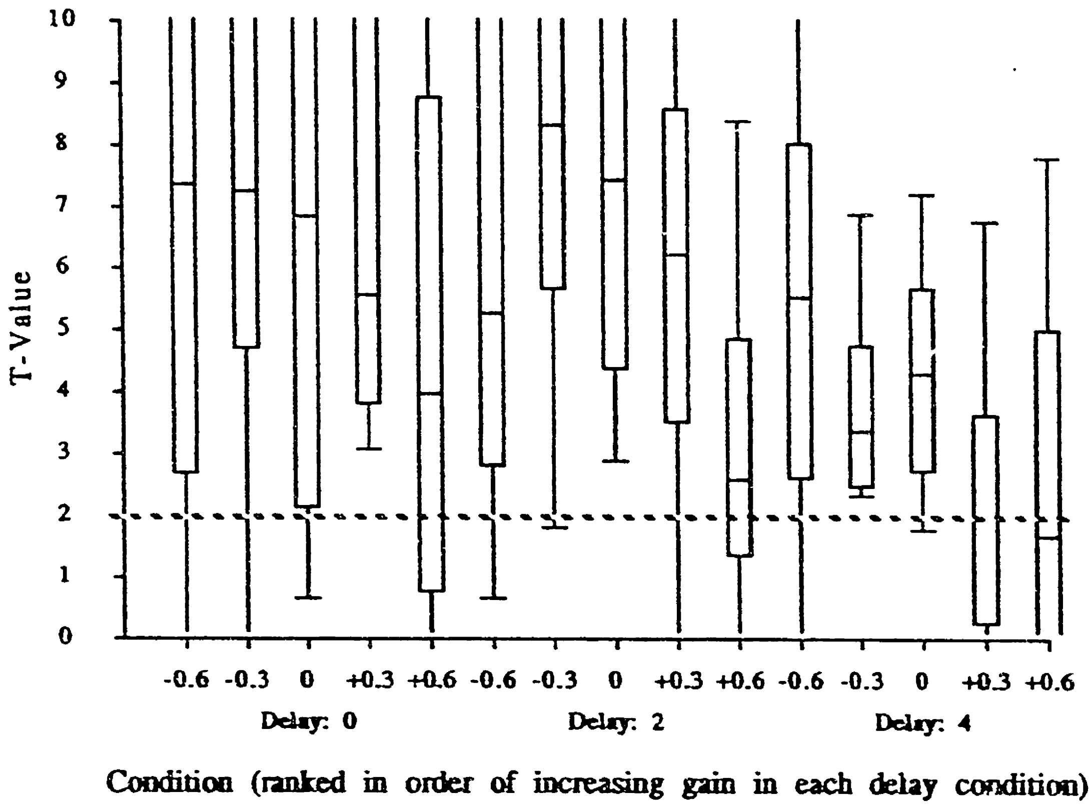

Figure 5.32 T-values for (future production - sales) weights: across treatments

With regard to fit of the weights, Figure 5.32 shows the t-values across conditions. Fit is significant for easier conditions, while increasing delay and high positive gains cause fit to decrease. Subjects' apparently decrease their importance of future change in inventory when the task becomes more difficult, unlike the optimal rule which increases attention given to future changes in inventory as difficulty increases.

Practice effects

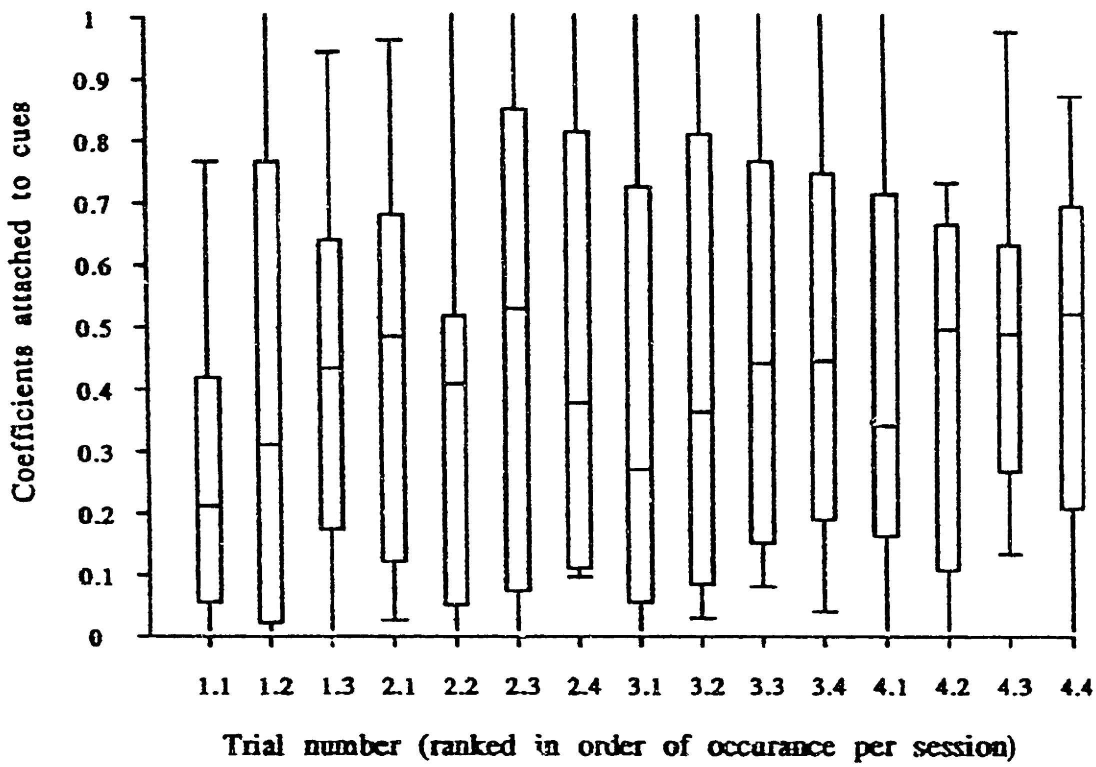

Figure 5.33 Variance of (future production - sales) weights: across trials

Effects due to the sequence of trials are blurred somewhat by the variability, but there appears to be a slight increase in weight given to future change in inventory after the first two trials, as Figure 5.33 shows. This effect is corroborated by the notebook analysis where it was seen that subjects' attention to change in inventory increased through the first two trials and leveled off thereafter. Most of the learning that occurs during the first two trials can be attributed to the increased weights attached to future change in inventory.

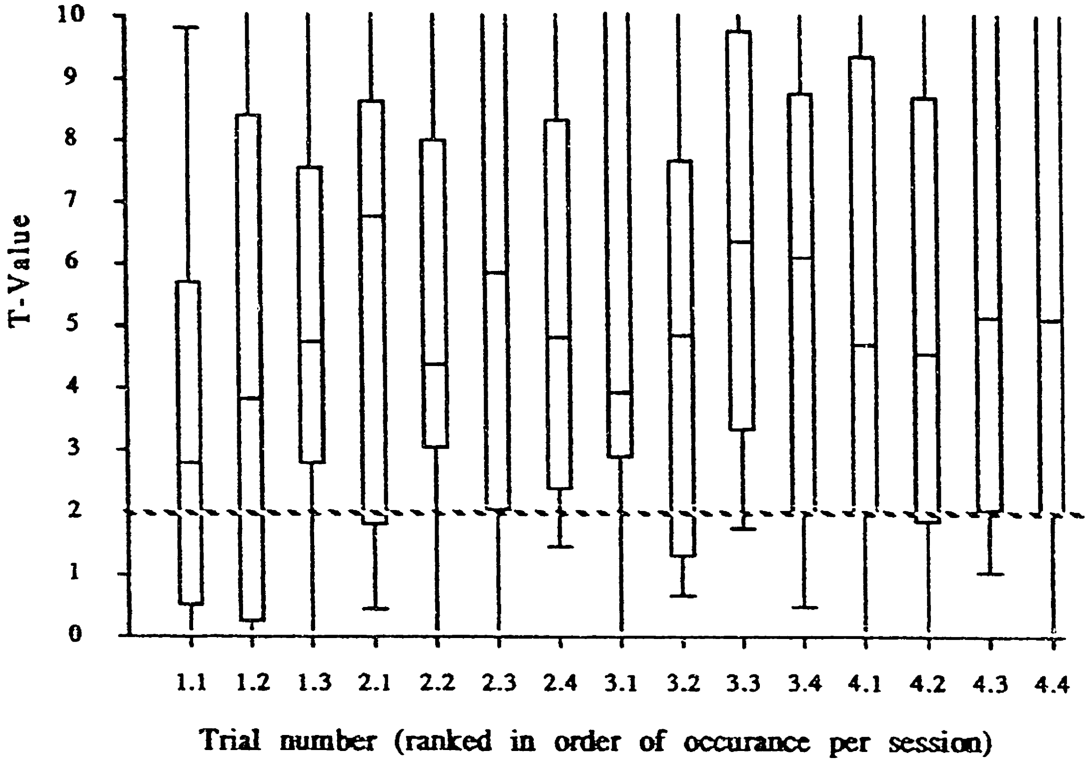

Figure 5.34 T-values for (future production - sales) weights: across trials

Figure 5.34 shows the same pattern in t-values, increasing fit after the first two trials and not much change in later trials. Although the variability is large, the slight increase at the beginning of the experiment and the consistency afterward indicate that the weights given to future change in inventory do not change much throughout the experiment.

Subject effects

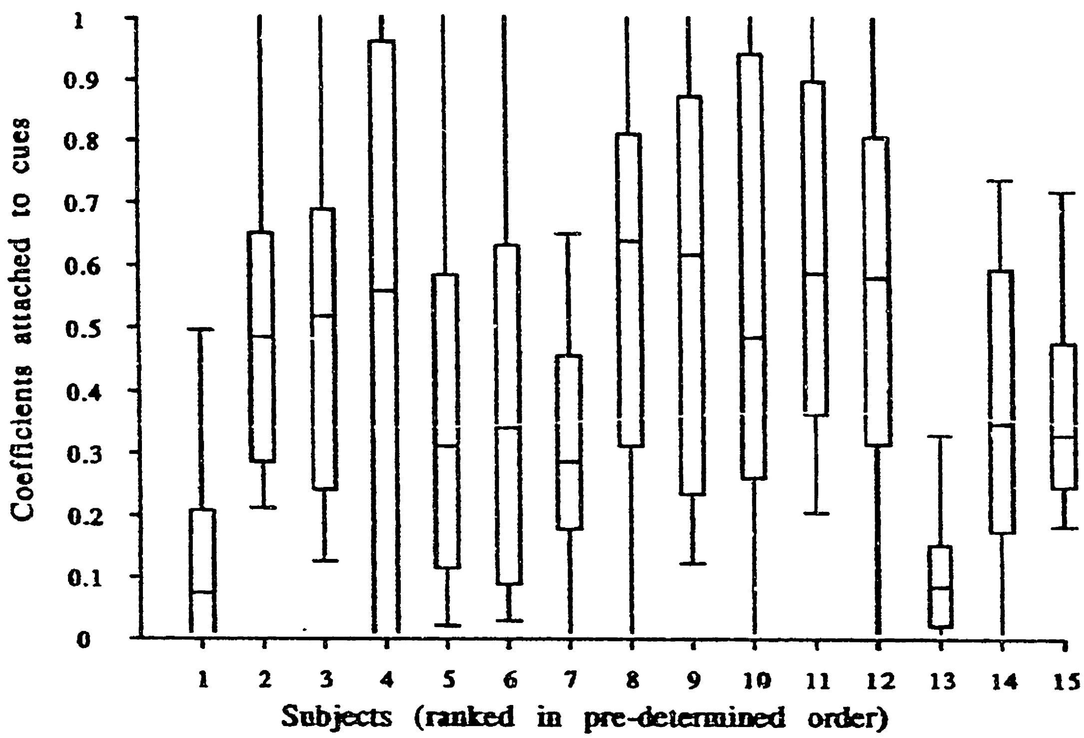

Figure 5.35 Variance for (future production - sales) weights: across subjects

Figure 5.35 shows that subjects do differ with respect to the attention paid to future change in inventory. Subjects 1 and 13 put very little attention on the supply-line which, in part, explains their low performance. The other subjects' median weights range from .3 to .6, which are more consistent with the optimum rule but are still less than optimum. The manner in which subjects adapt their weights as conditions change varies for different subjects. Subjects 13 and 15 remain within a very close range, while other subjects, such as subject 4, have a quite wide range of adjustments throughout the experiment. Thus, there appear to be some differences in the willingness of subjects to experiment with alternative weights.

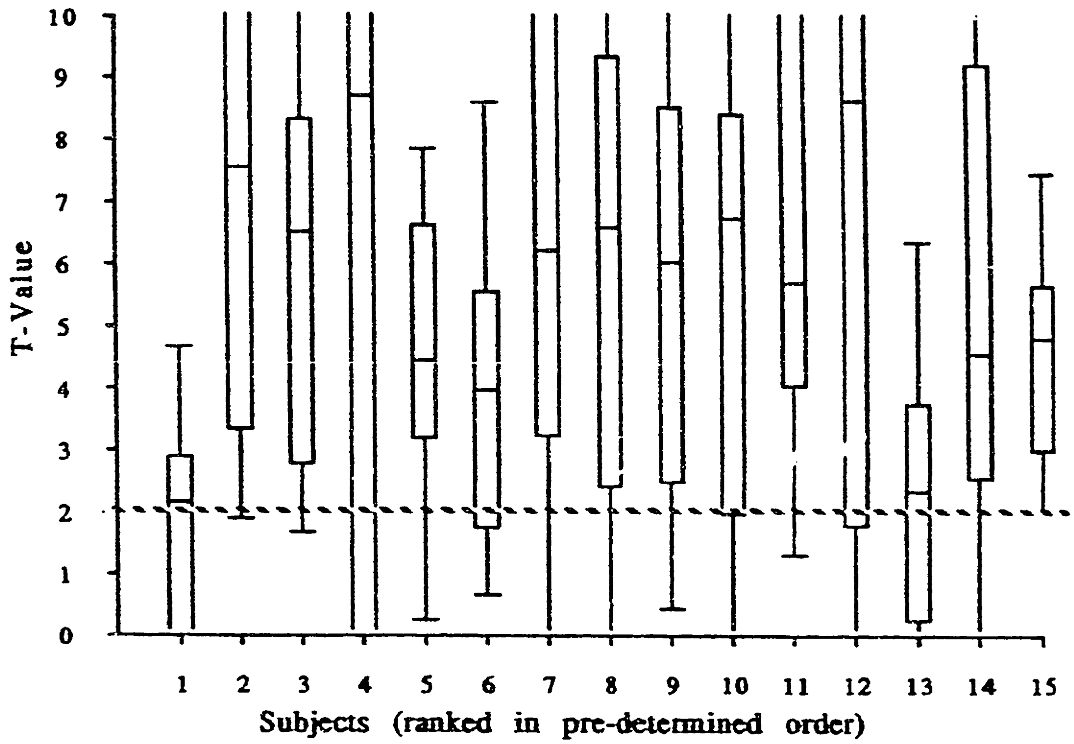

Figure 5.36 T-values for (future production - sales) weights: across subjects

Figure 5.36 shows the t-values for the subjects. T-values are consistent with the weight analysis: subjects 1 and 13 show decreased fit; barely half of the weights are significant. Subject 4's inconsistency is revealed again by large variability in the t-values, despite the high median t-value. All other subjects show significant fit The implications of the subject effects are that subjects differ with respect to the attention that they place on future changes in inventory. Low performers and high performers can be distinguished by the emphasis given to the supply line, which was also found in the notebook analysis.

5.3. Causes of Underperformance

We have identified a two cue rule that subjects employ and have analyzed the weights that they attach to those cues. On average, compared to optimum, subjects' scores are higher by a factor of four. What remains is an explanation for the discrepancy between optimum and actual behavior.

Underperformance can be the result of three principal causes: (1) Subjects consistently employ the wrong decision model; (2) Subjects employ the right model but apply it inconsistently or (3) Subjects employ the wrong model and apply it inconsistently. To understand the deep causes for dysfunctional decision making, we need to be able to distinguish to what extent underperformance is caused by employing the wrong model and to what extent underperformance is attributable to inconsistencies. We will see that the greatest cause for subjects' performance is a faulty mental model, indicated mainly by insufficient weights.

Figure 5.37 shows these relationships graphically and the different sources of variation that each comparison explains.

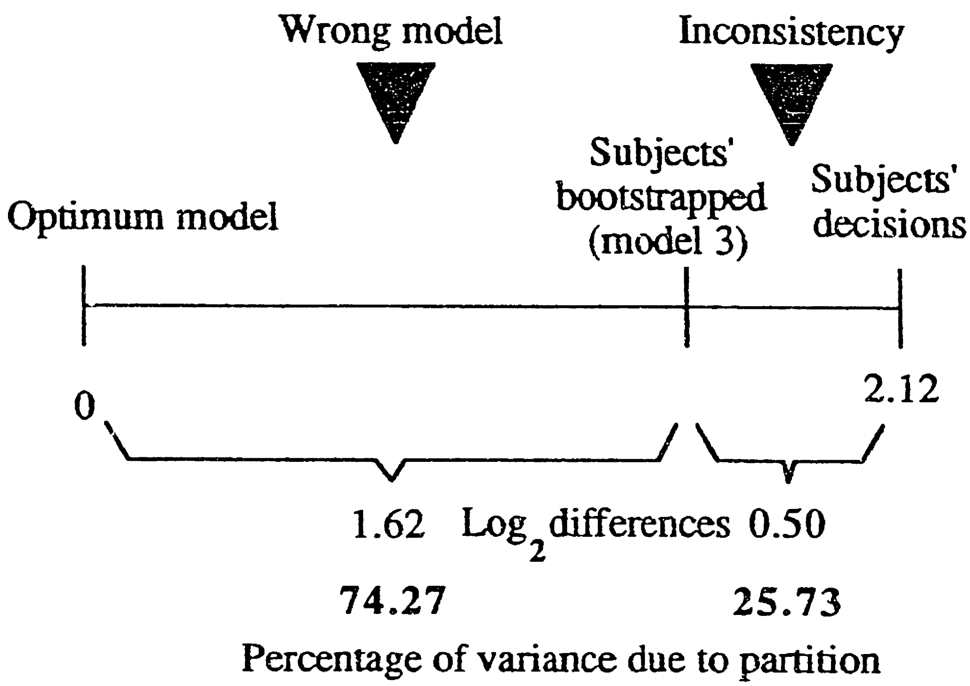

Figure 5.37 Partitioning subjects' variance as a function of $\log_2$ difference from optimum and subjects' bootstrapped rule: wrong model versus noise

In the following we will describe the results in more detail and explain the techniques used.

Bootstrap

We will use the technique of 'bootstrapping' to determine to what extent the observed underperformance can be contributed to random error in the subjects' judgment. Bowman (1963) was the first to demonstrate in a business context that management's own average behavior might yield better results than the behavior itself. He reasoned that bootstrapping works because people respond to irrelevant cues in the environment (like the last telephone call). Decision rules incorporating coefficients derived from management's own recurrent behavior eliminate these random and particularistic components. As Dawes (1971, p. 182) states,

> "A mathematical model, by its very nature, is an abstraction of the process itself; hence, if the decision maker's behavior involves following valid principles but following them poorly, these valid principles will be abstracted by the model - as long as the deviations from these principles are not systematically related to the variables the decision maker is considering."

To determine the bootstrapping effect we compare two models: subjects' actual performance and subjects' two-cue rule. Using the weights we estimated for each of the 225 trials, we rerun the experiment with the same random conditions for sales but substituting the estimated rule for subjects' decisions.

Model 3: Decision= a1*Inventory + a2*(Future Production-Sales)

Analysis

Since we were expecting that the results of the bootstrapping would differ across treatments and that the variance would be partially explained by practice and subject effects, we choose to use the standard toolkit, based on the ANOVA of the various effects.

Overall performance

Figure 5.38 shows the average bootstrap effect in $\log_2$ differences of costs between subjects' actual performance and the subjects' bootstrapped rule for the 15 conditions.

Overall, the estimated model outperforms subjects in most conditions; the exceptions can be seen wherever a negative difference is found. Negative differences imply that subjects outperform the estimated rule in those conditions. The bootstrap effect is most pronounced in the no-delay condition. Based on the regression analysis, we are not surprised to find smaller differences in the more difficult conditions. As we can see from ANOVA Figure 5.49, the treatment effects are significant and mainly attributable to differences in delay.

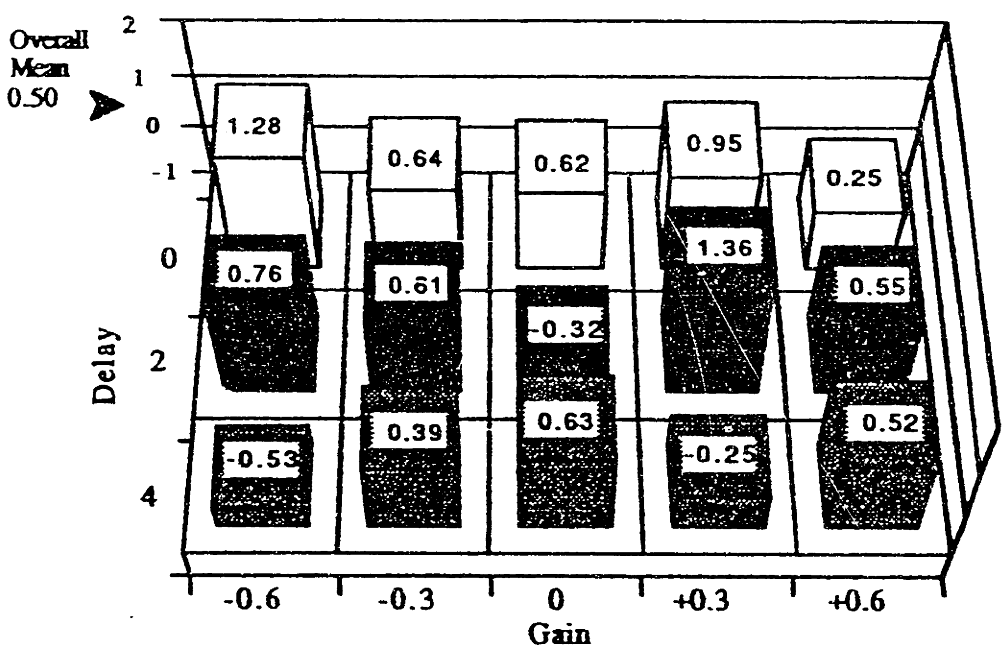

Figure 5.38 Bootstrap effects: $\log_2$ differences of costs between subjects' actual performance and subjects' bootstrapped rule

The overall mean of .50 $\log_2$ differences, indicating that the score for the bootstrapped model is roughly 30%

lower than subjects' scores. In some conditions, we see $\log_2$ differences of greater than 1, implying that subjects had over two times greater costs than the 2-cue, bootstrapped rule. Thus, there is a rather large impact of noise.

It seems important to note that a large bootstrap effect exists despite the fact that we have less confidence that the 2-cue model accurately describes subjects behavior under increasing delay.

Analysis

Source	df	F	probability
Overall treatment	14,182	1.09	$p<.05$
Delay	2,182	1.80	$p<.05$
Gain	4,182	0.22	------
Delay x Gain	8,182	1.35	------
Subject	14,182	1.20	------
Practice	14,182	1.52	------

-- not significant

Figure 5.39 Bootstrap ANOVA summary: subjects' decisions vs. estimated 2-cue subjects' rule

Treatment effects

ANOVA Figure 5.39 reveals that there is a treatment effect of delay ($p<.05$). Figure 5.40 provides a more detailed view of the bootstrap effect across the fifteen conditions by depicting the variance of $\log_2$ differences across conditions.

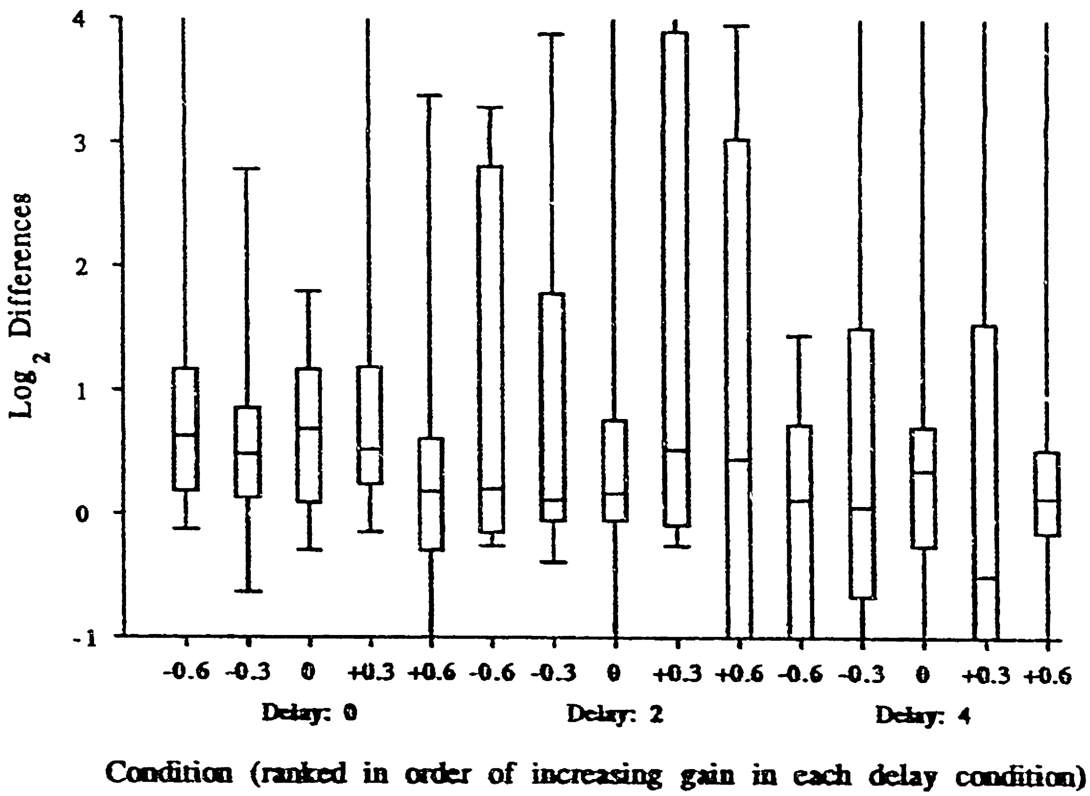

Figure 5.40 Bootstrap effects: $\log_2$ differences of costs between subjects' actual performance and subjects' bootstrapped rule- treatment effects

Higher bootstrap effect exists in the 0 and 2 delay conditions. In the 4 delay and 2 delay, +0.6 gain cases substantial number of cases show reverse bootstrap effect in which bootstrapping decreases performance rather than increasing it. We take this as another indication that model

3 might not capture subjects' rules completely, as we would have expected from the notebook analysis where we saw a large breakdown in heuristics under high difficulty.

Practice effects

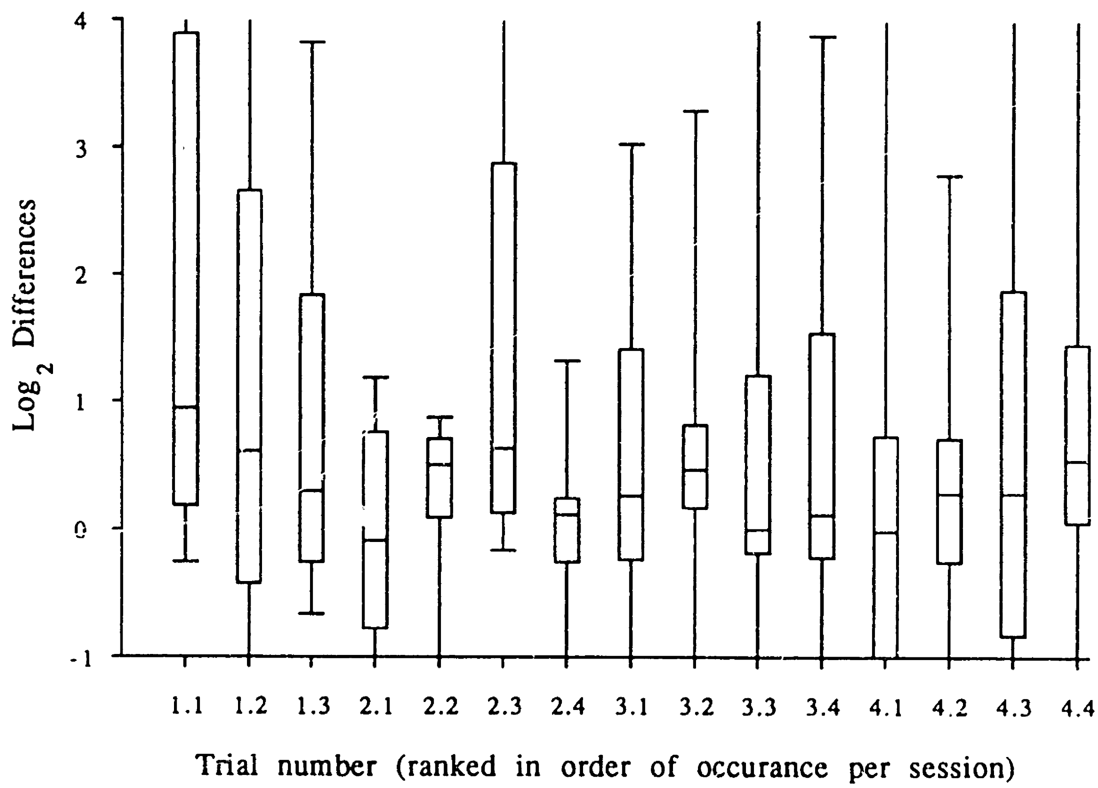

Figure 5.41 Bootstrap effects: log_2 differences of costs between subjects' actual performance and subjects' bootstrapped rule-practice effects

Although the practice effect is insignificant, it appears from Figure 5.41 that the bootstrap effect is highest during the first two trials. This result would confirm expectations. It can be assumed that inconsistency is highest in a phase when subjects are still in the process of familiarizing themselves with the task. During the first two trials, inconsistency might be attributable to deliberate experimentation.

Subject effects

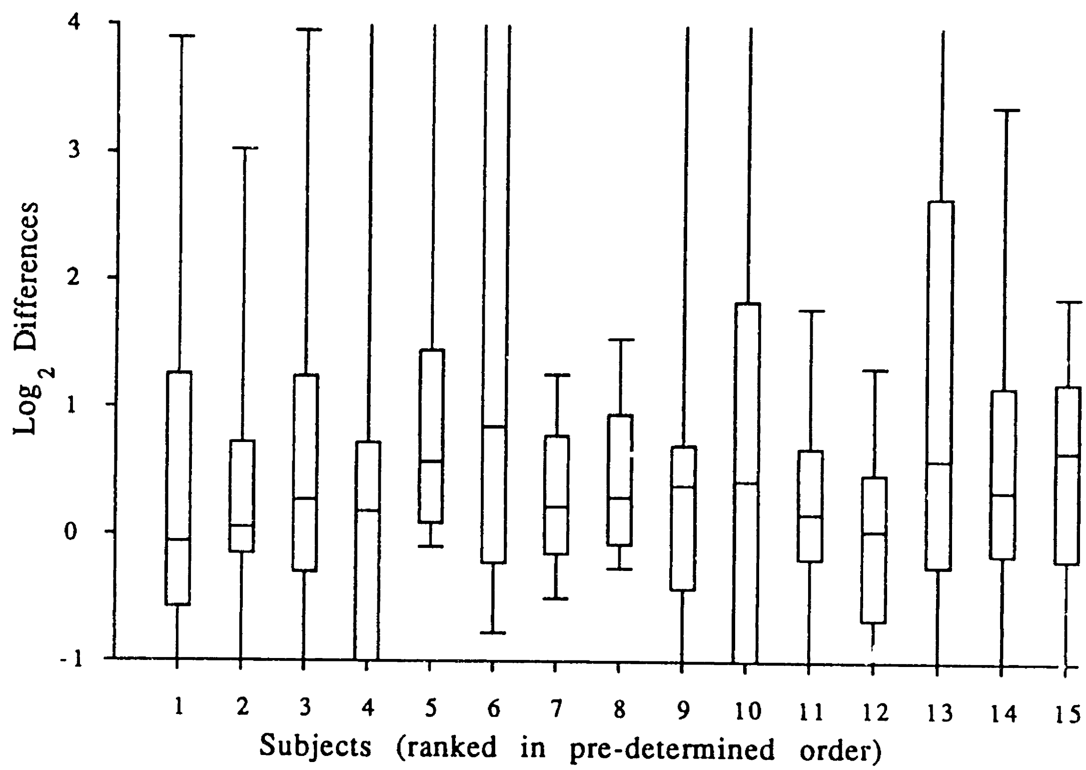

Figure 5.42 Bootstrap effects: $\log_2$ differences of costs between subjects' actual performance and subjects' bootstrapped rule- subject effects

Figure 5.42 shows the detailed effects for the 15 subjects; the ANOVA revealed no subject effects. Still, we can draw some conclusions from the figure. Subject 12's performance for more than half of the cases shows a reverse bootstrap effect, indicating that his mental model is somewhat more sophisticated than model 3 alone. Subjects 4, 10, and 13 on the other hand, do not appear to adhere to the two cue rule consistently, as corroborated by the notebooks.

Wrong model

In order to shed more light on the nature of subjects' mental models, it might be useful to distinguish between the effects of choosing certain cues versus placing improper emphasis on the cues considered.

Model 3 is not the full information model; how much of a difference in performance is accounted for by not utilizing the full state model vs. using model 3? To answer this, we compare the true optimum and the optimum weights for a two cue model. The costs for both rules are calculated through 225 simulation runs under the same conditions as the experiment.

The weights for the two-cue optimum were derived via grid search. In order to find the minimum costs per condition, each condition is simulated repeatedly with the weight on inventory varying systematically from 0 to 1 with a grain size of 0.01, and the weight on future change in inventory varied from 0 to -6 with a grain size of 0.01. To determine the optimum weights for a 2-cue model, the system is initialized with an inventory imbalance; no exogenous inputs were used. [1] Upon reaching equilibrium, inventory costs were recorded. At the end of the simulation runs, the

---

[1] See also page 282 for a more detailed look at the results of the grid search.

weights corresponding to the minimum costs were recorded as they are displayed in figure 5.43.

Delay \ Gain	-0.6	-0.3	0	+0.3	+0.6
0	a1=0.32 a2=0.49	0.35 0.58	0.39 0.71	0.43 0.92	0.48 1.32
2	0.32 0.74	0.33 0.94	0.34 1.27	0.35 1.89	0.35 3.38
4	0.28 0.85	0.27 1.11	0.25 1.50	0.24 2.28	0.22 3.97

Figure 5.43 Optimal weights of 2-cue "optimum" model

The weights derived were then used to simulate the experiment once more. Using the weights, we conducted a simulation experiment very similar to the one that we used to determine optimum scores. The only difference being that our 15 computer subjects would use only inventory and (future production - sales) in computing a decision.

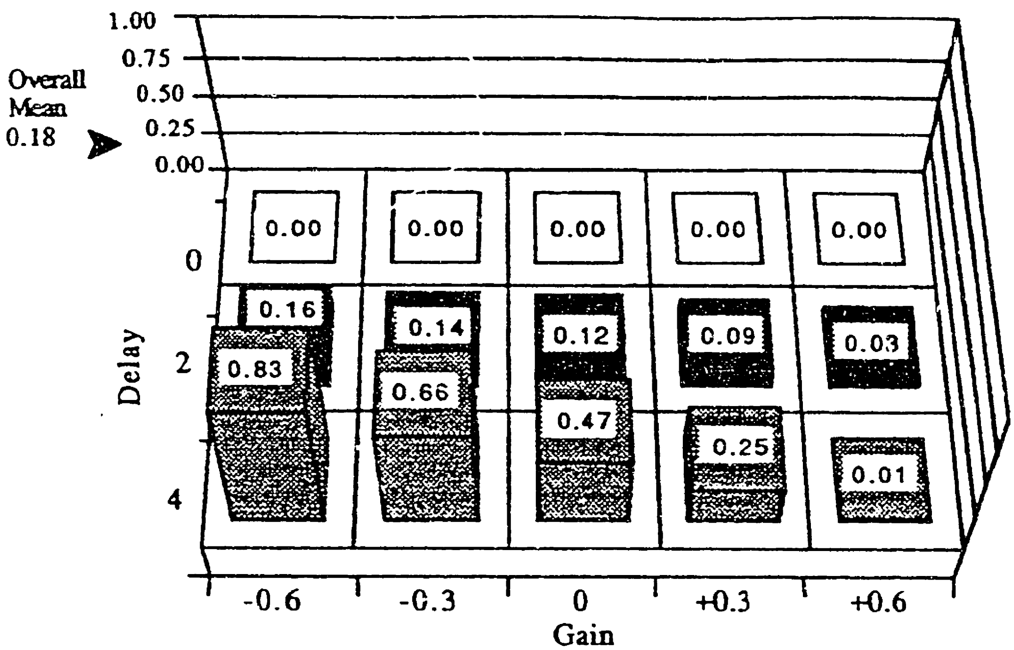

Figure 5.44 Optimal model vs. 2 cue optimum model: average $\log_2$ difference in costs per condition

Figure 5.44 shows the average $\log_2$ difference in costs between the optimum rule and an optimum 2-cue rule. The two models are identical for the no delay case, thus the difference is 0. We did expect the difference to increase as delay increased as more and more cues are ignored by the two-cue model. It is surprising, though, to find an apparent gain and delay interaction. The consequences for an insufficient cue model are higher under negative than under positive gain conditions. Overall, the effect of ignoring cues seems to be moderate.

In the first phase of the analysis, we calculated the variance accounted for by using insufficient cues. We may now calculate the effects of using insufficient weights by taking the $\log_2$ differences between the performance of an optimum 2-cue rule and subjects' bootstrapped rule. The resulting $\log_2$ score will reflect the variance associated with utilizing incorrect weights, and requires a more detailed examination, which is provided in the section that follows.

Figure 5.45a shows the $\log_2$ differences of variance accounted for by insufficient cues, weights, and noise. The overall difference between optimum and subjects' performance was determined to be roughly four times. Examination of Figure 5.45 a reveals that the solid black bars, signifying the difference between 2-cue optimum and the bootstrap subjects' rule, accounts for most of the variance in performance. The second largest contribution in variance arises from the variance caused by noise. Noise plays a more prominent role in the easier conditions, in part, because by definition there is no difference in optimum versus optimum 2-cue rule for 0 delay. The third and smallest contributor shows the variance accounted for by insufficient cues, which has relatively little influence. The negative components represent those cases when subjects' actual rules outperform the bootstrap rule which results in a negative $\log_2$ difference between subjects' and bootstrap rules.

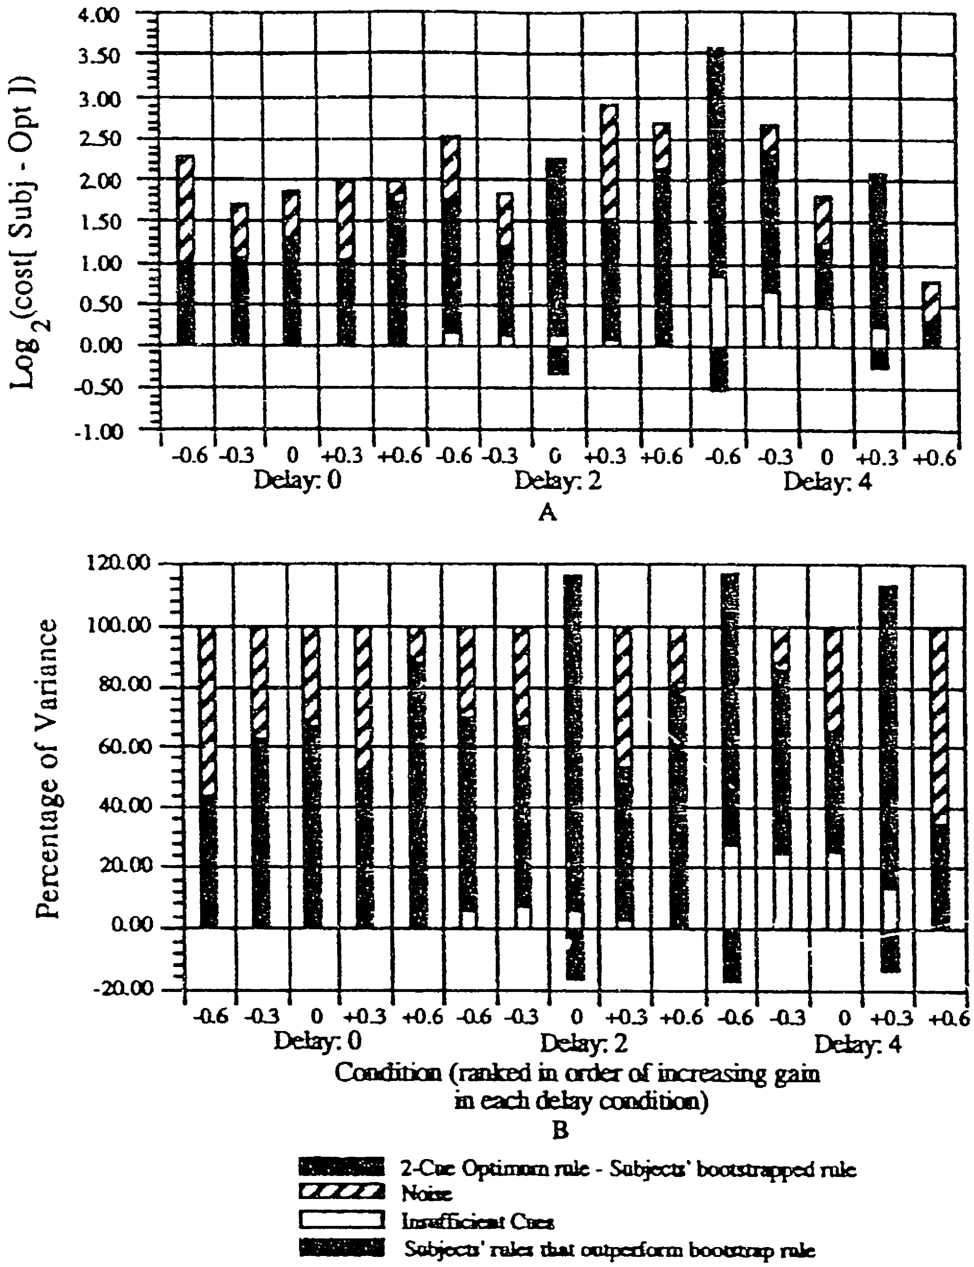

Figure 5.45 a & b Log_2 differences & percentages of variance: 2-cue optimum vs. 2-cue subjects' rules

Figure 5.45b shows the three variance components normalized to 100%. Noise accounts for 26% of the difference between subjects and optimum. Insufficient cues account for approximately 8% of performance, while insufficient weights account for roughly 66% of underperformance. By far, the fact that subjects hold incorrect models accounts for most of their underperformance.

We will now examine more closely the underperformance caused by insufficient weights. Figure 5.46 shows 15 vignettes arranged in four dimensions. Weight adjustment for subjects and optimal 2-cue rules are shown in the standard 3x5 matrix of conditions in which rows show the 3 delay conditions, and columns show the 5 gain conditions. Within each vignette, inventory weight adjustment is shown along the "y" axis ranging from 0 to 1, and future change in inventory is shown along the "x" axis, ranging from 0 to 2. Note that the units of inventory and expected change in inventory are the same, (since expected change in inventory is per period where period = 1), so the magnitudes of the cue weights for inventory and change in inventory are comparable. Subjects' weights are shown as "x's", and optimum weights are shown as a "+".

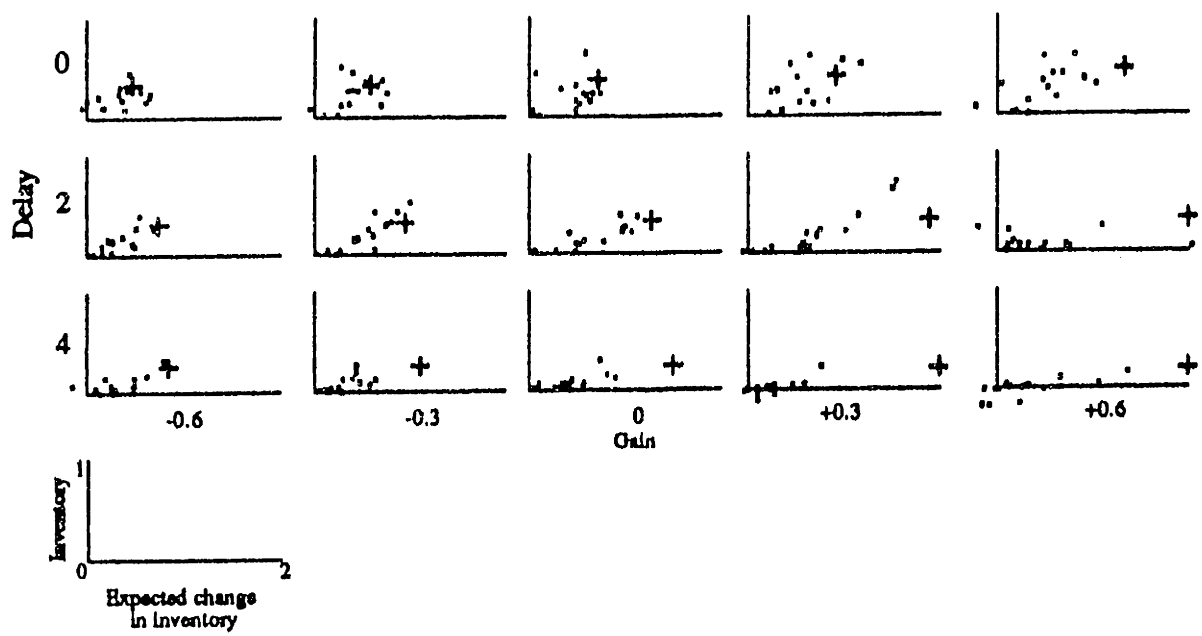

Figure 5.46 2-Cue optimum vs. 2-cue subjects' rule:
weight adjustment across conditions

Comparing subjects and optimum weight adjustment across conditions reveals that nearly all of the subjects' weights lie to the lower left of optimum, corroborating the earlier result of significant undercontrol, reformulated as insufficient weights applied to cues. The figure reveals two major points:

> 1. Under increasing complexity, the optimum weight for inventory stays roughly the same, but subjects' weights for inventory decrease dramatically.
>
> 2. Under increasing complexity, the optimum rule places greater and greater weight on future change in inventory, but subjects' weights remain the same.

It appears from inspection that the distance between the subjects' model and the optimum model increases with complexity. In order to verify these findings, an ANOVA was performed on the Euclidean distance of subjects' weights from optimum weights for inventory and future change in inventory, as depicted in figure 5.42. The Euclidean distance is calculated with the following formula for each subjects' weights:

$$\text{Distance from optimum weight} = \sqrt[2]{(\text{inventory weight})^2 + (\text{expected change in inv. weight})^2}$$

Overall performance

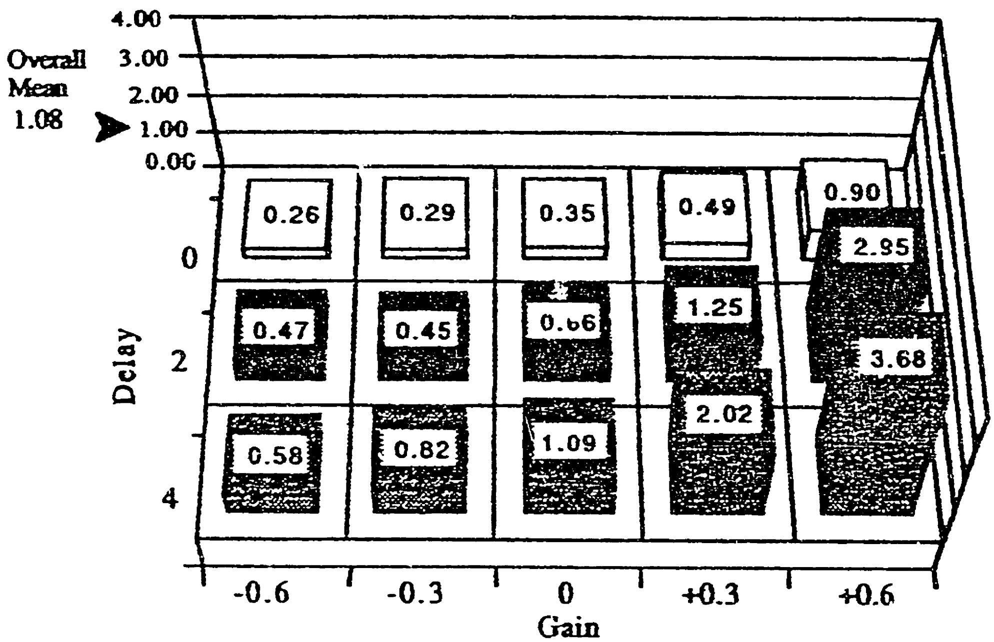

Figure 5.47 Euclidean distance of subjects' weights from optimum weights applied to inventory and future change in inventory - across conditions

Figure 5.47 shows that the Euclidean distance from optimum increases as difficulty increases. The average distance from optimum weights is 1.08. Since most of the subjects' weights fell to the lower left of optimum weights in Figure 5.46, then we may interpret the increasing distances from optimum weights as an additional indication that subjects undercontrol the system, especially under increasing task difficulty.

Analysis

Source	df	F	probability
Overall treatment	14,182	214.61	p<.01
Delay	2,182	361.25	p<.01
Gain	4,182	454.07	p<.01
Delay x Gain	8,182	58.22	p<.01
Subject	14,182	5.83	p<.01
Practice	14,182	0.78	------

-- not significant

Figure 5.48 ANOVA for Euclidean distance of subjects' weights from optimum for inventory and future change in inventory

ANOVA Figure 5.48 shows main effects of delay and gain as well as an interaction (p<.01). In addition, subject differences are significant (p<.01), but practice effects are not significant.

Treatment effects

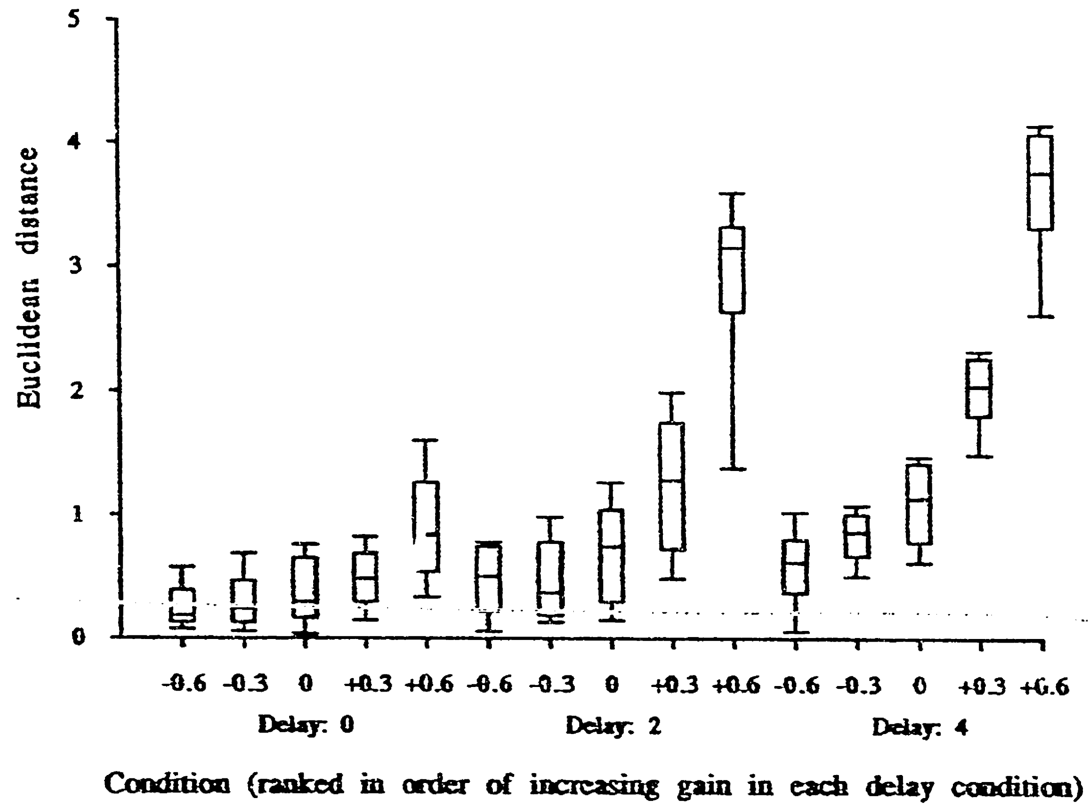

Figure 5.49 Euclidean distance of subjects' weights from optimum weights applied to inventory and future change in inventory- treatment effects

From figure 5.49 we see clear increases in distance from optimum weights as conditions increase in difficulty. As the ANOVA revealed, treatment effects of delay, gain, and their interaction are apparent. The interaction effect shown as steeper rise in distance across treatments reveals that as difficulty increases, subjects' weights move further away from the optimum weights, indicating greater and greater undercontrol on the part of subjects.

Practice effects

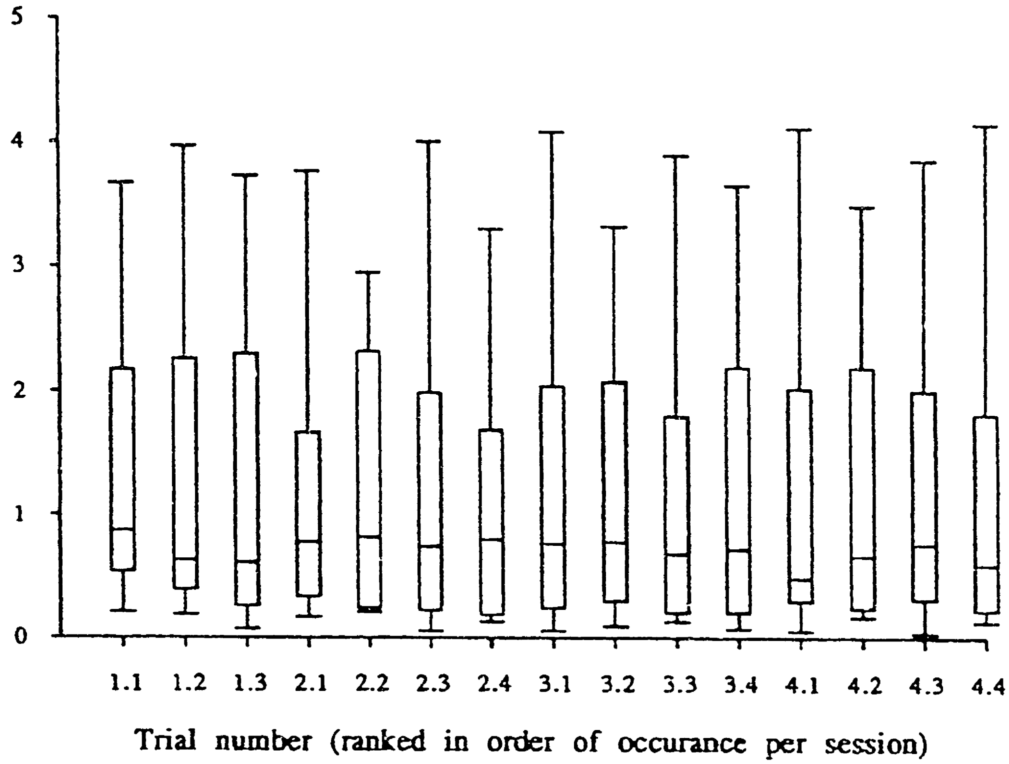

Figure 5.50 Euclidean distance of subjects' weights from optimum weights applied to inventory and future change in inventory- practice effects

Figure 5.50 shows that Euclidean distances remain fairly constant across trials as ANOVA figure 5.48 showed.

Subject effects

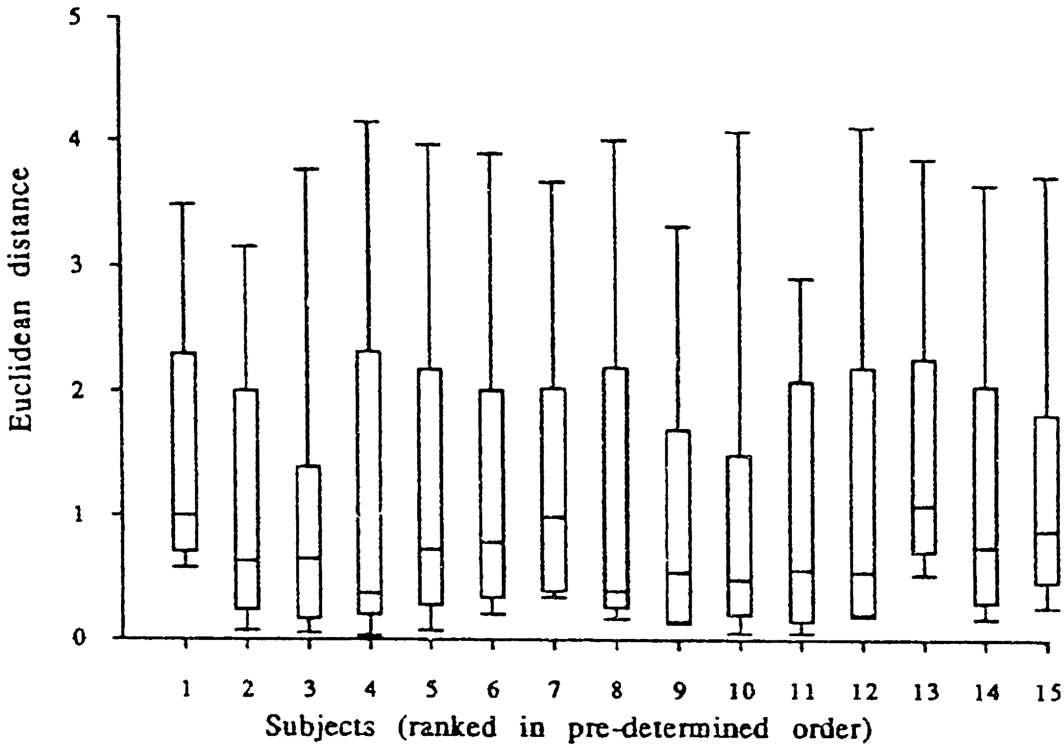

Figure 5.51 Euclidean distance of subjects' weights from optimum weights applied to inventory and future change in inventory- subject effects

Figure 5.51 confirms that the best performers, such as S11 and S12, in the majority of their trials stay closer to optimum than low performers, such as S1 and S13.

Given the vast discrepancies between optimum and subject weights, we are surprised that there were not greater increases in subjects' costs versus optimum. Refer to Figure 5.52 which shows the optimum weights for the fifteen conditions. Fifteen vignettes are arranged according to conditions, reflecting 3 delay and 5 gain cases. Within each vignette, the x-axis reflects the weight associated with the cue of future change in inventory, while the y-axis reflects the costs associated with the weight given to inventory. The scale on the x-axis ranges from 0 to 2, while the scale on the y-axis ranges from 0 to 1.

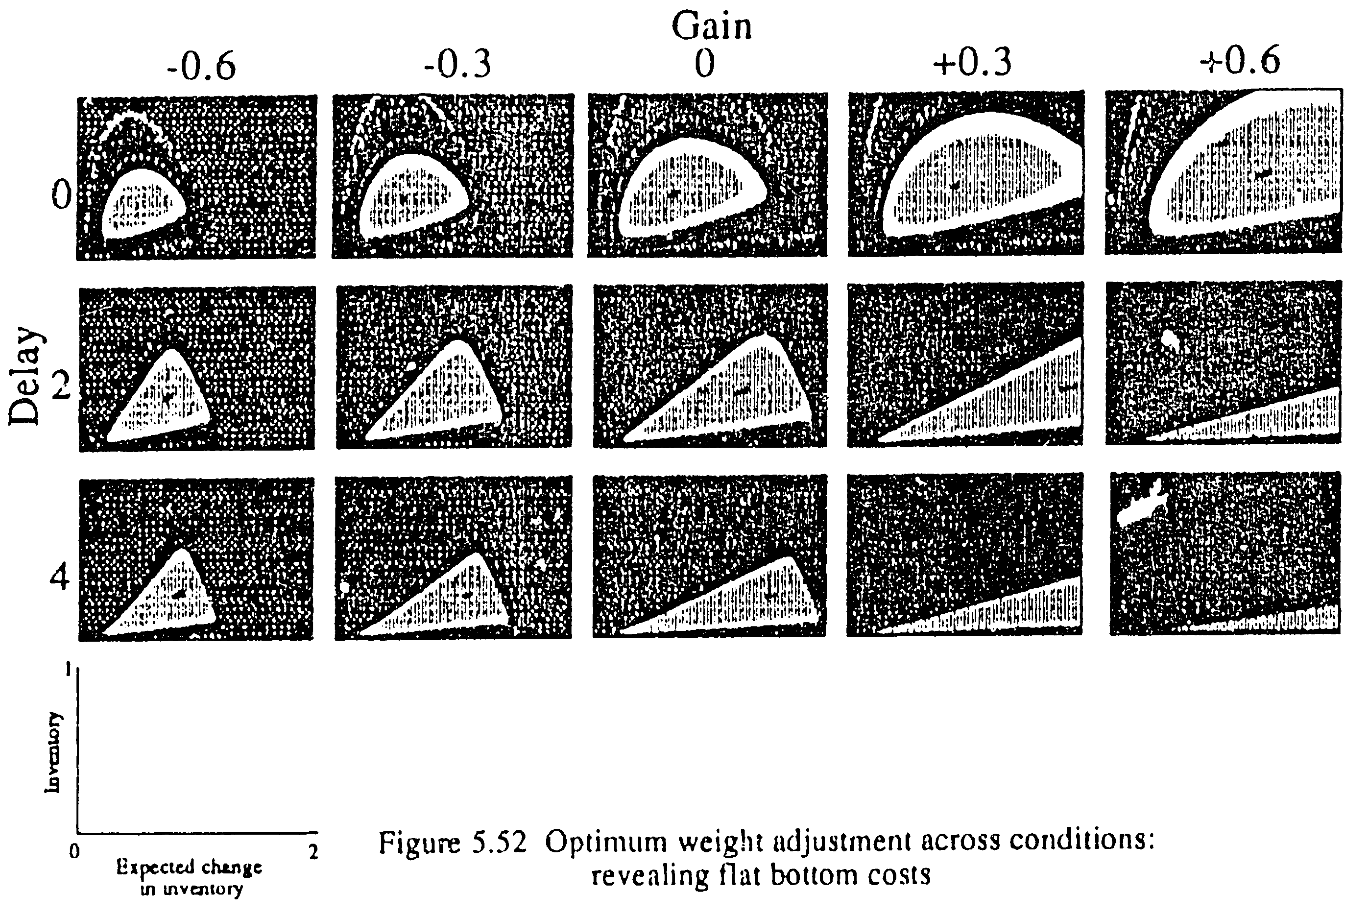

Figure 5.52 Optimum weight adjustment across conditions: revealing flat bottom costs

The black spot in the center of each vignette represents minimum costs. While the minimum costs achieved vary widely with the objective difficulty of the task, we have normalized the figures in such a way that each successive band around the center symbolizes a 50% increase in costs. Note that for conditions 2 delay, +0.6 gain, and 4 delay, +0.3 and +0.6 gain, the minimum cost is off the scale.

Notice that the slope of the objective difficulty does not increase linearly from the optimum performance to decreasing performance, rather, the slope of objective difficulty increases very slowly at first and then very rapidly. The area of low degradation is relatively large, and it just so happens that as difficulty increases, the area of low degradation narrows and stretches out toward the (0,0) point representing the strategy of "no-control," thus explaining that despite the huge differences in undercontrol, we do not see huge differences in costs. [2]

[2] A similar phenomenon has been reported in earlier studies (Rapoport, 1975). While not at the center of this thesis, the technique used in figure 5.52, showing systematically how the performance valley changes with delay and gain should provide impetus for further examination of the flat optimum phenomenon.

6. Conclusions

Summary of Analyses

The findings of the research presented thus far may be summarized as follows:

Benchmark analysis results

1. Objective difficulty of the task increases quite dramatically as measured by the optimum rule performance.

2. As objective difficulty of the task increases, controllability of the task decreases, as measured by the comparison of the no-control and the optimum rule.

3. Experimental research needs to consider that benchmark behavior may differ across treatment conditions rather than to assume that it is invariant with treatments.

Score analysis results

4. Subjects' raw scores increase dramatically with increasing task difficulty.

5. Subjects perform better than the no-control rule in most conditions but perform worse in the most difficult conditions.

6. Absolute differences between optimum and subject scores increase with increasing difficulty, while the ratio of subjects' to optimum scores remains roughly the same across conditions.

Time analysis results

6. Subjects spend the same amount of time in all conditions, despite large variations in task difficulty and number of available cues.

7. Subjects spend less time in the later trials.

8. Subjects differ greatly with respect to the amount of time spent, yet performance and time spent on decisions are not correlated.

9. Time spent on initial decisions is only slightly larger than time spent on later decisions within trials.

Control Effort results

10. Costs resulting from changes in production are far less than indicated by optimum, thus indicating undercontrol as further corroborated in later sections.

11. Undercontrol is correlated with score. High-performers exert less undercontrol than low-performers.

Notebook analysis results

12. Subjects spend little time in an orientation phase.

13. Most of the learning, measured as improvement in performance compared to optimum, occurs within the first two games for most of the subjects.

14. Increasing attention given to the change in inventory accounts for most of the performance improvement that occurs within the first three games.

15. Only some of the subjects show evidence of supply line control.

17. Subjects do not tend to vary their heuristics as treatment conditions vary unless forced to in a crisis situation. They then revert to a no-control approach.

18. Subjects show some signs of attributing system performance to exogenous factors.

Regress.. analysis results

19. Subjects do not utilize all of the cues available to them, but rather rely upon a subset that includes only inventory and expected future change in inventory.

20. Subjects place less weight on both of the cues than suggested by optimum.

21. While the optimum rule increases the weight on future change in inventory as gain increases, subjects' weight remains constant.

22. While the optimum rule places a constant weight on inventory as delay increases, subjects reduce the weight given to inventory.

23. Attention given to weights is the single most important factor in explaining the difference in performance between high and low performers.

24. Attention given to future states increases slightly as trials progress.

25. Faulty mental models account for 75% of subjects' performance, while inconsistency accounts for 25%.

26. Insufficient weight given to the two cues accounts for the greatest portion of the difference between subjects' and optimum performance: roughly 66%.

27. Despite large differences in the weights attached to cues, subjects do not show larger differences in scores compared to optimum, which is the result of "flat bottom" characteristics of the task, that allows for a wide range of near-optimal performance.

We can condense our findings into two major conclusions:

> 1. People maintain the same limited-cue adjustment heuristics in response to situations with varying complexity. These heuristics break down under increasing task difficulty.
>
> 2. People learn the limited-cue adjustment heuristic quickly and are reluctant to modify it.

These conclusions will now be discussed in light of the current study's findings and previous research.

1. People maintain the same limited-cue, adjustment heuristics in response to situations with varying complexity. These heuristics break down under increasing task difficulty.

We see substantial misperceptions in subjects' mental models with respect to delay and gain. Based on the preceding findings, we hold the hypothesis that people do not use a full information model, but rather employ a simple, two-cue, adjustment heuristic, as evidenced by the time, control effort, notebook, regression, and weights analyses.

Three main explanations for the development of subjects' heuristics have been offered in the literature.

> a. People consciously make a cost-benefit trade-off of limited cognitive resources (Payne, 1982).
>
> b. People generate and decide between alternative decision rules, which they test by experimentation (Wallsten, 1980).
>
> c. People rely upon faulty mental models that do not capture the dynamic nature of the task. In particular, people misattribute endogenous system behavior to exogenous factors (Sterman, 1989b).

a.) Previous research has cited the mental effort required to perform the task as a possible cause for the development of subjects' heuristics (Payne, 1982). While subjects may have the mental resources to perform in quite complex tasks, it may be that people make a necessary cost-benefit trade-off of cognitive effort (Kleinmuntz, 1990):

> In the stock management task, there are a number of factors that may increase the effort required to incorporate the supply line into the decision. For instance, while the current sales figure and the current inventory level are directly visible to the decision maker, the supply line may have to be reconstructed from memory (e.g. recalling the previous order quantities and aggregating those numbers to arrive at an overall estimate). If the supply line is harder to use than the other information, the decision maker may decide to do without it. If the costs associated with decreased performance are not severe, then this may lead to a satisfactory compromise: A lot of effort is saved and little is sacrificed.

However, in this task, supply line information is always available to subjects and is presented in the same display as other cues such as inventory and sales. In this task failure to utilize supply line information cannot be attributed to costs of acquiring the information.

Brehmer (1989), in addition, pointed out that people do not treat all delays as equal, but rather have greater difficulty with those components of the system that must be inferred from system structure. In our task, all information was made transparent. Nonetheless, subjects performed worse under increasing delay and gain. The severe underperformance and limited-cue models suggest that subjects cannot mentally utilize a rule of any great complexity, (findings 10, 14, 15, 17, 19, 20, 21, 22, 23, 25). Reduction of mental effort alone seems not to be a sufficient explanation for subjects' performance, since there is a possibility that subjects' limited-cue models may be caused by their inability to generate competing hypotheses.

b.) One strategy for performing the task is to generate multiple decision rules and decide, based on score, in favor of the most promising rule. However, in our study, people did not seem to generate competing hypotheses to perform the task, but rather maintained one rule and modified it only when it failed. Kleinmuntz (1990) likewise proposed that people may not entertain alternative decision rules, but choose to use an initial rule until it fails. Once the rule fails, the person may modify it by trial and error in order to develop a new, more effective rule.

We find that high performers developed and maintained a heuristic that includes influences of delay quickly. Low performers generate poorer quality heuristics, which do not attend to delay, early and must abandon them sooner than high performers. Even high performers' heuristics failed under high delay. Thus, we saw both groups of subjects develop one rule and utilize it until it failed. Although our subjects did show signs of modifying their rules upon breakdown, the majority of subjects' heuristics break down and are substituted by an emergency "hands off" approach, rather than generating alternative hypotheses. Once the crisis situation is over, subjects tend to revert to their original heuristics only slightly modified (findings 13, 14, 15, 17, 19). There seems to be a difference between heuristic breakdown in crisis and deliberate generation of alternative decision strategies. Previous researchers, (Kleinmuntz, 1990; Sterman, 1989b), have investigated the generation of subjects' initial rule.

c.) Sterman's (1989a) research revealed several misperceptions of feedback. Sterman proposed that subjects' heuristics are locally rational, implying that their decision rules did not capture the dynamic relationships within the task:

> If the decision rule is locally rational, the explanation for the poor performance of the subjects must be sought in the interaction between the decision rule and the feedback structure of the simulated economy.

Our consistent finding of lack of attention to the supply line seems to confirm Sterman's hypothesis. While rules that ignore the supply line are consistent with no-delay situations, they become dysfunctional once delay or strong positive feedback side effects are introduced (findings 14, 15, 17, 19, 21, 24, 26).

Kleinmuntz (1990) agrees with Sterman that subjects' initial rule may be locally rational:

> "An interesting question is how the initial rule is generated. The rule may be one that is "locally rational" in the sense that it is a very general strategy that provides effective performance in simple static tasks that resemble the dynamic task at hand. For example, the anchor and adjustment rule is very effective in stock management problems that do not have lagged feedback structures." (Sterman, 1989b, pp.323-324)

From this perspective, the problem of which rule subjects decide upon becomes a framing issue. What mental models do subjects generate in response to the training and task characteristics?

The notebooks and regression analyses reveal that subjects' rules are fairly established by the second trial. There are no dramatic changes in the number of cues considered. For instance, there is a predictable progression from inventory to change in inventory to future change in inventory within the first two to three trials for most subjects. Under moderate task difficulty, 0-2 delay and negative gain, subjects' heuristics perform fairly well, but under increased difficulty, high delay and high gain, subjects' heuristics fail. In addition we found that the weights that subjects place on all cues are less than indicated by the optimum rule, leading to underperformance.

Sterman's subjects misattribute system performance to exogenous factors. Likewise, our subjects seem to be incapable of generating appropriate closed-loop models, as evidenced by ignoring vast parts of the supply line, and our subjects did show signs of attributing system behavior to exogenous factors (finding 18). At the same time, however, we hesitate to label subjects' models as open- loop models. The fact that subjects are inclined to take a hands-off approach in the most difficult decisions might very well indicate that subjects are acutely aware both that they are part of the system and that their insufficient understanding of the system could be a cause of system instability. Subjects might want to close the loop but do not know how to do so effectively.

2. *People learn the limited-cue adjustment heuristic quickly and are reluctant to modify it.*

The following section will explore the processes underlying the development and persistence of the two-cue hypothesis.

Most learning occurs in the first two trials. The two-cue model fits subjects' data fairly well under most conditions, with high delay and high gain being exceptions. Subjects' heuristics were developed early in the experiment and only modified, if at all, when those heuristics broke down. We did find changes in strategy during those games. We think that we can attribute those changes neither to deliberate generation of alternative strategies as discussed in paragraph 1) above, nor to thematic vagabonding (Dörner, Kreuzig, Reither, and Stäudel, 1983), as argued below.

Dörner, et al.'s (1983) work with LOHHAUSEN, which was described in the literature review, showed that the subjects exhibited "thematic vagabonding," or generating new hypotheses without testing them appropriately. High performers tested one causal hypothesis completely, while low performers tended to jump from one hypothesis to another. In delay situations, much time can pass between control actions and system responses. Disproving incorrect hypotheses can only be achieved by waiting for the feedback of those control actions; most people do not incorporate these delays into their search for appropriate control policies. Contrary to Dörner, we found that few subjects were willing to experiment with varied control strategies, and subjects generally held onto their initial strategies until they proved ineffective to the task, perhaps due to difficulty of benchmarking one's self.

It seems to us that we can characterize the observed changes in strategy best as a temporary abandonment of initial strategy. The human factors literature similarly reports that people revert to previously learned responses in emergency situations (Singleton, 1978).

An important corollary to heuristic change is the amount of effort spent in processing those changes. Interestingly, more effort spent does not correspond with increased learning. Time data show tnat time spent on the task and scores achieved are not correlated. From the control effort results, we know that subjects undercontrol compared to optimum. From the weights analysis, we learned that subjects' performance could mainly be attributed to insufficient weights given to certain cues. Thus, subjects' heuristics show a faulty development and little improvement over trials, despite instances of increased time spent per decisions. Decreased performance despite increased decision times suggests three hypotheses:

> 1) The extra time is spent in irrelevant thought, calculation, rest, or distraction.
>
> 2) People vary in intrinsic processing speed independent of intelligence in this task. Since optimal is flat, low effort is not much penalized in this task.
>
> 3) The people who understand the task less think harder compensating for their poor understanding.

The fact that few subjects experiment with the system despite ample opportunity and that heuristics are developed quickly and do not change much might be taken as an indication that the learning process was terminated prematurely. It appears that a useful representation of the task including a full understanding of the consequences of differing values for delay and gain is not reached.

The first major conclusion discussed subjects as controllers and the second discussed subjects as learners. To fully understand dysfunction in decision making, we need to understand the inherent trade-off of control and learning as it is influenced by task complexity.

Subjects face the challenge of controlling the system while learning about it at the same time. Task complexity affects both processes: the ease with which the system can be controlled and the ease with which the underlying relations within the system can be learned. To explore the relationship between control and learning we propose a framework as illustrated in Figure 6.1. We can identify three major loops: 1. a control loop, 2. a passive learning loop, and 3. an active learning loop.

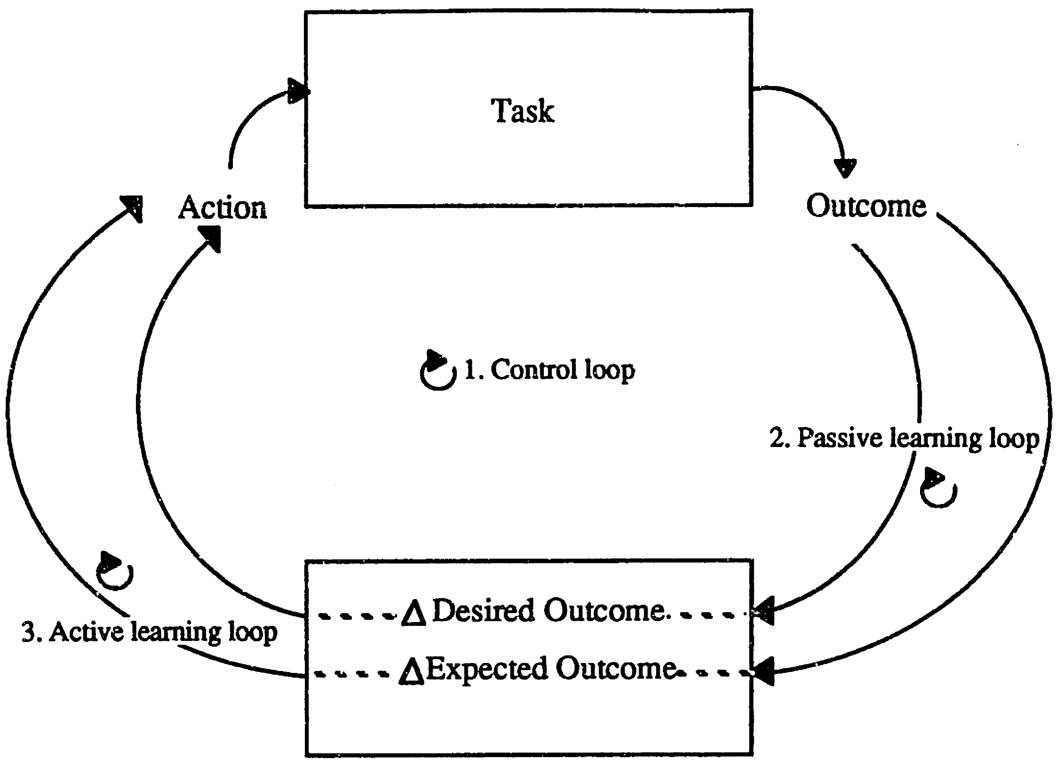

Figure 6.1 Learning and control loops contained in a dynamic decision task:

1. Control loop
2. Passive learning loop
3. Active learning loop

In this framework the distinction between learning and control depends on the way in which outcome feed-in enters the decision process. Comparing outcomes with expected outcomes is central to learning, while comparing outcomes with desired outcomes is central to control.

In order to control the system, the decision maker needs only to compare selected outcomes to desired outcomes. In the task at hand the desired outcome for inventory is 0; any discrepancy between actual inventory and desired inventory is acted upon depending upon the decision maker's current understanding and mental model of how his actions can reduce the goal discrepancy. Understanding of the system develops, apart from the a-priori instructions received, either through the passive or active learning loops.

The two learning loops differ in the degree of intervention the decision maker uses to produce the outcomes. In the active learning loop, the decision maker chooses actions to test assumptions he might have about the system. This is very much in line with Rasmussen (1986). Some of decision maker's tests might be extreme condition tests in which the task, as such, is simplified through testing its response to extreme inputs (Rasmussen, 1986). The passive learning loop does not contain learning motivated actions, but rather relies solely upon control motivated actions, which are initiated to reduce goal discrepancies. In the passive learning mode, learning is a by-product of control. These two different motivations comprise the two links of action feed-out.

Having developed this framework, we are now in a position to discuss how both active-loop learning and passive loop learning are affected by the need for control and how task complexity affects the trade-off.

Active learning loop

Even in situations with the lowest task complexity, we did not find many signs of active-loop learning. We are not surprised, since even in low complexity situations, learning-motivated actions may increase, rather than decrease, goal discrepancies. Active-loop learning may be effective in the long run but is costly in the short term. More often than not learning-motivated actions are discouraged in real systems, and understandably so. Should a nuclear reactor melt down or a plane crash, people would not be comforted by knowing that the operator was simply attempting to learn the system by experimenting with alternative control actions.

We did find some limited examples of active learning. As observed in the notebook analysis, one subject attempted a risk-free experiment by making production decisions in the hundredths and thousandths decimals, so as to produce consequences in hundredths and thousandths decimals. This innovative attempt lasted for only a few rounds of play, suggesting that even experimentation does not lead to learning, since the outcomes were to hard to understand.

One interpretation of the persistent finding of undercontrol is that subjects never become confident enough with their mental models to act decisively. Since the system deters the use of the active learning loop through the consequences of increased costs, subjects may not reach a confidence level necessary for good control which would be an indication of active generation of alternative control strategies. Related to this speculation is the idea that although people possess knowledge, they are unwilling to use it until a certain subjective criterion confidence level is reached (Hunt, 1989). Thus, subjects may have hypotheses regarding a better model than is suggested by the notebook, regression, and weight analyses, but they are simply unwilling to test them for fear of driving costs too high by making a "mistake." The active learning loop, therefore, becomes secondary to the goal of keeping costs low. In a sense, learning in the experiment is determined by subjects' willingness to experiment despite consequences.

Passive learning loop

What we label "passive learning" in this study is very similar to the concept of outcome-feedback as generally used in the behavioral decision literature. (In the following we will use the term outcome "feed-**in**" to refer to this concept in order to reserve the label "feed-**back**" to circumstances where a loop is closed.) The similarity becomes especially apparent in the most difficult task conditions where action feed-out is dramatically reduced and subjects turn almost into observers of an open-loop system. We know from static outcome feed-in experiments that outcome feed-in is not very effective (Brehmer, 1980). Since demands on memory become more stringent in a dynamic outcome feed-in experiment, because measurements cannot be taken at one point in time but must be obtained repeatedly, we should expect outcome feed-in to become even less effective. The addition of control-motivated action feed-out does not seem to alter the conclusion: Outcome feed-in stays ineffective even when it is part of an outcome feedback **loop**.

People, although provided with complete information on structure and parameter relationships, do not completely infer the dynamic nature of the system and do not know what to expect. More precisely, people try but err in inferring dynamics from structure, and people do not use all the information and fall back on prior notions. Or, as Hogarth (1987, p.121) argues:

> "...even when such information is available, people do not necessarily use it. In particular, they have a tendency to seek information that confirms existing notions rather than to seek information that could disconfirm their hypotheses."

We can interpret the fact that learning seemed to stop after the second to third trial as a sign that the passive learning loop lost its effectiveness at this point. Subjects mainly develop their heuristics, albeit limited-cue heuristics, within the first 2-3 trials, indicating that the training received before the experiment was successful in providing a tentative mental model with which to continue the learning process. After trial 3, however, learning virtually plateaus suggesting that subjects do not continue to increase their understanding of the system.

The ability to learn ultimately depends upon the capability to form expectations, be it pro-actively or retro-actively. It could be that the main difference between high-performers and low-performers is their differing ability to generate satisfying expectations that can be tested. As delay and gain increase, low-performers reach an earlier point where their expectations fail and where they seem to be unable to replace failed expectations with better ones.

Both the active and the passive learning loop are at their weakest when they are needed the most: in difficult situations. Decision makers in dynamic tasks seem to be caught in a dilemma: The need to control overrides the ability to learn.

Implications

How, then can we equip the decision maker to better cope with complex dynamic tasks? How can we design systems that strengthen at least one of the learning loops, if not both?

In many real systems, the penalties for error are great and built-in punishments discourage experimentation and active learning. Decision makers may become conditioned over time to use only passive learning and only to the extent necessary to produce a sufficiently good heuristic. Moreover, the dynamics of actual decisions are considerably slower than in a simulation and decision makers may not be able to observe, much less infer the long-term consequences of decisions. These factors, combined with the bias to attribute bad events to exogenous factors, severely limit the effectiveness of both learning loops.

The creation of "practice fields" for professionals has been advocated as a possible way out of the learning dilemma (Diehl, 1992; Schön, 1983; Senge, 1990). Practice fields are artificial "microworlds" that share many task characteristics with the real "macroworld". At the same time, they are designed to eliminate all those obstacles to learning that are typically present in the real world. Typically, (1) actions are reversible; (2) time is speeded up to allow the experience of consequences that otherwise might not show for months or years; (3) mistakes are not punished but seen as possible learning opportunities; and (4) complexity is simplified to allow better understanding of the decision dynamics.

While newer developments in software technology have considerably eased the development of sophisticated microworlds (Diehl, 1990), research on the effectiveness of those practice fields has only started. We still do not have the answers to such essential questions as: (1) How transferable to the macroworld are insights gained within the microworld? (2) Do microworlds provide for a more effective learning than traditional methods such as seminars or books? and (3) What pedagogical theories and design principles should guide the development and use of microworlds?

As long as we do not have the answers to these basic questions it is too early to tell if microworlds can live up to the expectations of their proponents: helping decision makers learn both how to learn and how to control.

REFERENCES

Atkinson, R. C., Herrnstein, R. J., Gardner, L., and Luce, R. D. (1988). Steven's Handbook of Experimental Psychology. New York: John Wiley & Sons, Inc.

Attwood, D. (1974). The interaction between human and automatic control. In E. Edwards, & F. P. Lees (Ed.), The human operator in process control (pp. 120-134). London: Taylor and Francis.

Bainbridge, L., Beishon, J., Hemming, J. H., and Splaine, M. (1974). A study of real-time human decision making using a plant simulator. In E. Edwards, & F. P. Lees (Ed.), The human operator in process control (pp. 91-104). London: Taylor and Francis.

Bakken, B. E. (1990). Learning and transfer in simulated dynamic decision environments: an information processing approach (Sytem Dynamics Group Working Paper). Massachusetts Institute of Technology.

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. Acta Psychologica, 44, 211-233.

Bergin, R. S., and Prusko, G. F. (1990). The learning laboratory. Healthcare Forum, 33(2), 32-36.

Best, J. B. (1986). Cognitive Psychology . New York: West Publishing Company.

Bowman, E. H. (1963). Consistency and optimality in managerial decision making. Management Science, 9(2), 310-321.

Brehmer, B. (1980). In one word: Not from experience. Acta Psychologica, 45, 223-241.

Brehmer, B. (1989). Feedback delays and control in complex dynamic systems. In P. M. Milling, & E. O. K. Zahn (Ed.), Computer-based management of complex systems (pp. 189-196). Berlin, Heidelberg: Springer.

Broadbent, D. E., and Aston, B. (1978). Human control of a simulated economic system. Ergonomics, 21(12), 1035-43.

Brunswick, E. (1952). The conceptual framework of psychology. Chicago: University of Chicago Press.

Carroll, J. S., and Johnson, E. J. (1990). Decision Research. Newbury Park: Sage Publications.

Carver, C. S., and Scheier, M. F. (1982). Control theory: A useful conceptual framework for personality-social, clinical, and health psychology. Psychological Bulletin, 92(1), 111-135.

Chamberlin, E. H. (1948). An experimental imperfect market. Journal for Political Economy, 56, 95.

Chapman, L. J. (1967). Illusory correlation in observational report. Journal of Verbal Learning and Verbal Behavior, 6, 151-155.

Coleman, J. S. (1986). Psychological structure and social structure in economic models. The Journal of Business, 59(4), S365-S377.

Crossman, E. R. F. W., and Cooke, J. E. (1974). Manual control of slow-response systems. In E. Edwards, & F. P. Lees (Ed.), The human operator in process control (pp. 53-66). London: Taylor and Francis.

D'Azzo, J. J., and Houpis, C. H. (1981). Linear Control System Analysis and Design: Conventional and Modern . New York: McGraw-Hill Book Company.

Dawes, R. M. (1971). A case study of graduate admissions: Application of three principles of human decision making. American Psychologist, 26, 180-188.

Diehl, E. W. (1988). Participatory simulations as training tools - A study based on the market growth model. In J. B. Homer, & A. Ford (Ed.), Proceedings of the 1988 international conference of the system dynamics society (pp. 52-64). La Jolla, CA: System Dynamics Society.

Diehl, E. W. (1989). A study on human control in stock-adjustment tasks. In P. M. MIlling, & E. O. K. Zahn (Ed.), Computer-based management of complex systems (pp. 205-212). Berlin, Heidelberg: Springer.

Diehl, E. W. (1990). MicroWorld Creator Software and User Manual . MicroWorlds Inc.: P.O. Box 1400, Kendall Square, Cambridge, MA, 02142.

Diehl, E. W. (1992). Participatory simulations for managers: The design philosophy behind MicroWorld Creator. Forthcoming in European Journal of Operations Research, 59(1).

Dörner, D. (1975). Wie Menschen eine Welt verbessern wollten.. Bild der Wissenschaft, 12(2), 48-53.

Dörner, D., Kreuzig, H. W., Reither, F., and Stäudel, T. (1983). Lohhausen. Vom Umgang mit Unbestimmtheit und Komplexität. Bern: Huber.

Dörner, D., and Reither, F. (1978). Über das Problemlösen in sehr komplexen Realitätsbereichen. Zeitschrift für Experimentelle und Angewandte Psychologie, 25, 527-551.

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

Edwards, E., and Lees, P., Frank. (1972). Man and computer in process control . London: The Institution of Chemical Engineers.

Edwards, E., and Lees, P. F. (1974). The human operator in process control . London: Taylor and Francis.

Edwards, W. (1962a). Dynamic Decision Theory and Probabilistic Information Processing. Human Factors, 4(2), 59-73.

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

Einhorn, H. J., and Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. Psychological Review, 85, 395-476.

Fleishman, E. A., Quaintance, M. A., and Broedling, L. A. (1984). Taxonomies of human performance: The description of human tasks . New York: Academic Press.

Forrester, J. W. (1980). Industrial Dynamics . Cambridge: The MIT Press.

Funke, J. (1988). Using simulation to study complex probelm solving: A review of studies in the FRG. Simulation & Games, 19(3), 277-303.

Garner, C. A. (1982). Experimental evidence on the rationality of intuitive forecasters. In V. L. Smith (Ed.), Research in Experimental Econonmics Volume 2 Greenwich: JAI Press.

Hogarth, R., M. (1981). Beyond Discrete Biases: Functional and Dysfunctional Aspects of Judgmental Heuristics. Psychological Bulletin, 90(2), 197-217.

Hogarth, R. M. (1987). Judgement and Choice (2 ed.). New York: John Wiley & Sons.

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

Holt. (1980). Competitive bidding for contracts under alternative auction procedures. J. Pol. Econ., 88, 433.

Huber, O. (1986). Mental Representation in Multistage Decision Making. ,, Bounded Rationality (pp. 66-81).

Hunt, D. (1989). Human self-assessment process theory: An eight factor model of human performance and learning and everyman's causation. In G. Ljunggren, & S. Dornic (Ed.), Psychophysics in Action (pp. 41-53). Berlin: Springer-Verlag.

Kampmann, C. E. (1989). Feedback complexity and market adjustment: An experimental approach (Sytem Dynamics Group Working Paper). Massachusetts Institute of Technology.

Kantowitz, B., H., and Sorkin, R., D. (1983). Human Factors: Understanding people-system relationships . New York: John Wiley & Sons.

Kleinmuntz, D. N. (1985). Cognitive Heuristics and Feedback in a Dynamic Decision Environment. Management Science, 31(6), 680-702.

Kleinmuntz, D. N. (1990). Information processing and the misperception of feedback in dynamic decision making. In D. F. Andersen, G. P. Richardson, & J. D. Sterman (Ed.), Proceedings of the 1990 International System Dynamics Conference (pp. 826-41). Chestnut Hill, Massachusetts: System Dynamics Society.

Kleinmuntz, D. N., and Thomas, J. B. (1987). The Value of Action and Inference in Dynamic Decision Making. Organizational Behavior and Human Decision Processes, 39, 341-364.

Kluwe, R. H., Misiak, C., and Haider, H. (1989). Modelling the process of complex system control. In P. M. MIlling, & E. O. K. Zahn (Ed.), Computer-based management of complex systems (pp. 335-342). Berlin, Heidelberg: Springer.

Kragt, H., and Landeweerd, J. A. (1974). Mental skills in process control. In E. Edwards, & F. P. Lees (Ed.), The human operator in process control (pp. 135-145). London: Taylor and Francis.

Kuipers, H. A., and Vennix, J. A. M. (1990). Organizing home care in the future: Using system dynamics to assess organizational changes. In D. F. Andersen, G. P. Richardson, & J. D. Sterman (Ed.), Proceedings of the 1990 International System Dynamics Conference (pp. 596-604). Chestnut Hill, Massachusetts: System Dynamics Society.

Langer, E. J. (1975). The illusion of control. Journal of Personality and Social Psychology, 32, 311-328.

Lees, F. P. (1974). Research on the process operator. In E. Edwards, & F. P. Lees (Ed.), The human operator in process control (pp. 386-425). London: Taylor and Francis.

Lees, F. P., and Sayers, B. (1976). The behaviour of process operators under emergency conditions. In T. B. Sheridan, & G. Johannsen (Ed.), Monitoring behavior and supervisory control (pp. 331-341). New York: Plenum Press.

Leventhal, H., Meyer, D., and Nerenz, D. (1980). The common sense representation of illness danger. In S. Rachman (Ed.), Medical Psychology New York: Pergamon.

Lucas, R. E. (1986). Adaptive behavior and economic theory. The Journal of Business, 59(4), S401-S426.

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

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

Mann, T. L., and Hammer, J. M. (1986). Analysis of user procedural compliance in controlling a simulated process. IEEE Transactions on Systems, Man and Cybernetics, 16, 505-510.

McLeod, P. (1976). Control strategies of novice and experienced controllers with a slow response system (a zero-energy nuclear reactor). In T. B. Sheridan, & G. Johannsen (Ed.), Monitoring behavior and supervisory control (pp. 351-358). New York: Plenum Press.

Meadows, D. (1989). Gaming to implement system dynamics models. In P. M. MIlling, & E. O. K. Zahn (Ed.), Computer-based management of complex systems (pp. 635-640). Berlin, Heidelberg: Springer.

Miller, G. A. (1956). Psychology's block of marble. Contemporary Psychology, 1.

Moray, N. a. L., and Pam Pajak, J. (1986). Acquisition of process control skills. IEEE Transactions on Systems, Man and Cybernetics, 16, 497-504.

Morecroft, J. D. W. (1988). Strategic microworlds and system dynamics modelling. In J. B. Homer, & A. Ford (Ed.), Proceedings of the 1988 international conference of the system dynamics society (pp. 282-297). La Jolla, CA: System Dynamics Society.

Morris, N. M., and Rouse, W. B. (1985). The effects of type of knowledge upon human problem solving in a process control task. IEEE Transactions on Systems, Man and Cybernetics, 15, 698-707.

Nemhauser, G. L. (1966). Introduction to dynamic programming. New York: Wiley.

Payne, J. W. (1982). Contingent decision behavior. Psychological Bulletin, 92, 382-402.

Peterson, C. R., and Beach, L. R. (1967). Man as an intuitive statistician. Psychological Bulletin, 68, 29-46.

Plott, C. R. (1982). Industrial organization theory and experimental economics. Journal of Economic Literature, 20, 1485-1527.

Plott, C. R., and Smith, V. L. (1978). An experimental examination of two exchange institutions. Rev. Econ. Stud., 45, 133.

Powers, W. T. (1973). Feedback: Beyond Behaviorism. Science, 179, 351-356.

Rapoport, A. (1966a). A Study of Human Control in a Stochastic Multistage Decision Task. Behavioral Science, 11, 18-32.

Rapoport, A. (1966b). A Study of a Multistage Decision Making Task with an Unknown Duration. Human Factors, 8(1), 54-61.

Rapoport, A. (1967). Variables Affecting Decisions in a Multistage Inventory Task. Behavioral Science, 12, 194-204.

Rapoport, A. (1975). Research paradigms for studying dynamic decision behavior. In D. Wendt, & C. A. J. Vlek (Ed.), Utility, Probability, and Human Decision Making (pp. 349-375). Dordrecht, Holland: Reidel.

Rapoport, A., and Wallsten, T., S. (1972). Individual Decision Behavior. Annual Review of Psychology, 23, 131-176.

Rasmussen, J. (1986). Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering . Amsterdam: North Holland.

Rasmussen, J., and Rouse, W. B. (1981). The human detection and diagnosis of system failures. New York: Plenum Press.

Reichart, U. (1986). Die Steuerung eines nichtlinearen Regelkreises durch Versuchspersonen. (Working Paper, Dept. of Psychology). University of Bamberg.

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

Richardson, G. P., Vennix, J. A. M., Andersen, D. F., Rohrbaugh, J., and Wallace, W. A. (1989). Eliciting group knowledge for model-building. In P. M. MIlling, & E. O. K. Zahn (Ed.), Computer-based management of complex systems (pp. 343-357). Berlin, Heidelberg: Springer.

Sanders, M. S., and McCormick, E. J. (1987). Human Factors in Engineering and Design (6 ed.). New York: McGraw-Hill Book Company.

Schön, D. (1983). The Reflective Practitioner: How Professionals Think in Action . New York: Basic Books.

Senge, P. M. (1990). The Fifth Discipline: The Art & Practice of the Learning Organization . New York: Doubleday.

Senge, P. M., and Lannon, C. (1990). Managerial microworlds. Technology Review, 93(5), 62-68.

Senge, P. M., and Sterman, J. D. (1990). Systems thinking and organizational learning: Acting locally and thinking globally in the organization of the future. In D. F. Andersen, G. P. Richardson, & J. D. Sterman (Ed.), Proceedings of the 1990 International System Dynamics Conference (pp. 1007-1022). Chestnut Hill, Massachusetts: System Dynamics Society.

Sheperd, A., Marshall, E. L., Turner, A., and Duncan, K. D. (1977). Diagnosis of plant failures from a control panel: A comparison of three training methods. Ergonomics.

Sheridan, T. B., and Ferrell, W. R. (1974). Man-machine systems: Information, control, and decision models of human performance . Cambridge, MA: MIT Press.

Sheridan, T. B., and Johannsen, G. (1976). Monitoring behavior and supervisory control . New York: Plenum Press.

Shubik, M. (1979). On the number of types of markets with trade in money: Theory and possible experimentation. In V. L. Smith (Ed.), Research in Experimental Economics Volume 1 Greenwich: JAI Press.

Shuford, E. H. (1964). Some Bayesian learning processes. In M. W. Shelly II, & G. L. Bryan (Ed.), Human judgements and optimality (pp. 127-152). New York: Wiley.

Simon, H. (1982). Models of Bounded Rationality . Cambridge, Mass.: MIT Press.

Simon, H. A. (1969). The science of the artificial . Cambridge, Mass.: MIT Press.

Singleton, W. T. (1978). The Study of Real Skills V1: The Analysis of Practical Skills . Baltimore: University Park Press.

Smith, Williams, Bratton, and Vannoni. (1982). Competetive market institutions: Double auctions vs. sealed bid-offer auctions. American Economic Review, 72, 58.

Smith, V. L. (1962). An experimental study of competitive market behavior. Journal for Political Economy, 70, 111.

Smith, V. L. (1965). Experimental auction amrkets and the Walresian hypothesis. Journal for Political Economy, 73, 387.

Smith, V. L. (1967). Experimental studies of discrimination versus competition in sealed-bid auctions. The Journal of Business, 40, 56.

Smith, V. L. (1976). Experimental economics: Induced value theory. American Economic Review, 66, 274.

Smith, V. L. (1982). Microeconomic systems as an experimental science. American Economic Review, 72(5), 923-955.

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

Sterman, J. D. (1989b). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. Management Science, 35(3), 321-339.

Tinker, R. F. (1990). Teaching theory building modeling: Instructional materials and software for theory building. In D. F. Andersen, G. P. Richardson, & J. D. Sterman (Ed.), Proceedings of the 1990 International System Dynamics Conference (pp. 1152-1163). Chestnut Hill, Massachusetts: System Dynamics Society.

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

Tversky, A., and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 5, 207-232.

Tversky, A., and Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. Science, 185, 1124-1131.

Tversky, A., and Kahneman, D. (1986). Rational choice and the framing of decision. The Journal of Business, 59(4), S251-S278.

Vennix, J. A. M., and Geurts, J. L. A. (1987). Communicating insights from complex simulation models. Simulation & Games, 18(3), 321-343.

Vennix, J. A. M., Gubbels, J. W., Post, D., and Poppen, H. J. (1988). A structured approach to knowledge acquisition in model development. In J. B. Homer, & A. Ford (Ed.), Proceedings of the 1988 international conference of the system dynamics society (pp. 419-433). La Jolla, CA: System Dynamics Society.

Vennix, J. A. M., and Scheper, W. J. (1990). Modeling as organizational learning: An empirical perspective. In D. F. Andersen, G. P. Richardson, & J. D. Sterman (Ed.), Proceedings of the 1990 International System Dynamics Conference (pp. 1199-1210). Chestnut Hill, Massachusetts: System Dynamics Society.

Wallsten, T. S. (1980). Cognitive Processes in Choice and Decision Behavior. Hillsdale, NJ: Erlbaum.

Walras, L. (1954). Elements of pure economics . Homewood, Ill.: Irwin.

West, B., and Clark, J. A. (1974). Operator interaction with a computer-controlled distillation column. In E. Edwards, & F. P. Lees (Ed.), The human operator in process control. London: Taylor and Francis.

Williams. (1973). The effect of market organization on competitive equilibrium: the multi-unit case. Rev. Econ. Stud. 40, 97.

Zeckhauser, R. (1986). Comments: Behavioral versus rational economics: What you see is what you conquer. The Journal of Business, 59(4), S435-S449.