EDUCATION

Teaching Takes Off

Flight Simulators for Management Education

By John D. Sterman

Chaos reigns in my classroom. Eighty students are shouting, gesturing and laughing while counting poker chips and turning over cards. A thick roll of $1 bills awaits the winners. A field trip to Las Vegas? No, it's the "Beer Game," a role-playing simulation designed to teach principles of management science.

We all know the world is growing more complex. Technological, social and environmental change are accelerating. Organizations, industries and government grow ever more tightly coupled. Today's students will face a world that is more dynamic and more uncertain than ever before.

Managers are not alone in facing such daunting tasks. Our society depends on systems of enormous complexity, from nuclear power plants to jumbo jets. Indeed, a popular metaphor likens managers and pilots. Managers must fly their organizations through uncharted skies and rough weather, constantly monitoring their information systems for signs of trouble or opportunity, dogfighting with the competition, preventing hijacking by hostile raiders – all the while giving the stockholders in the back a smooth ride.

There is one difference between managers and pilots, however. No airline

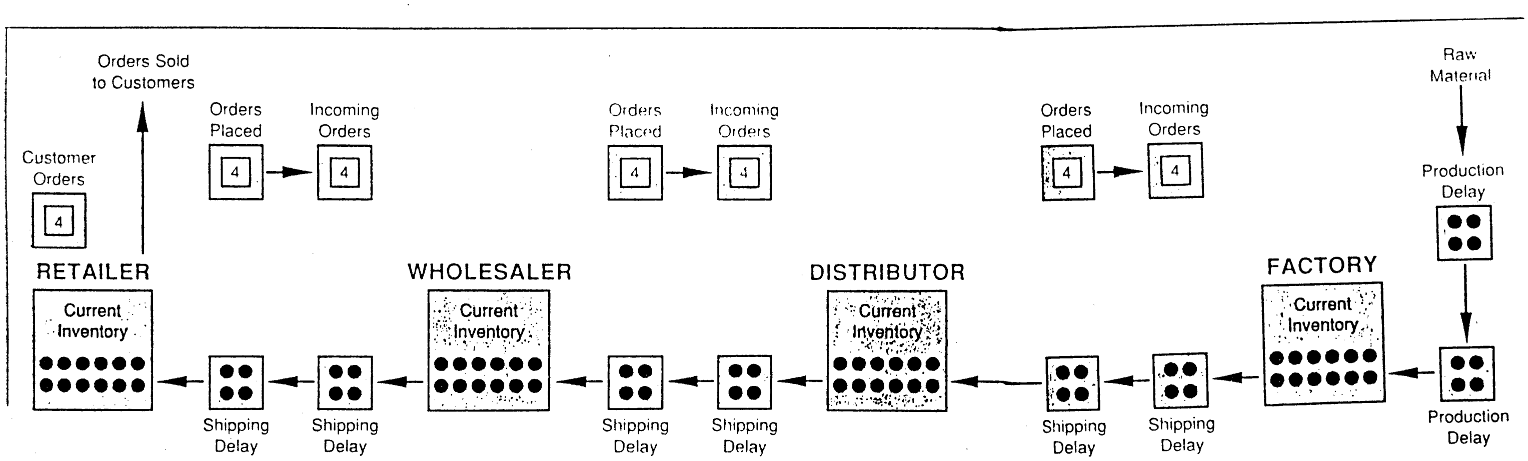

Figure 1: Beer Game board, showing initial conditions.

would dream of sending pilots up in a real jumbo jet without extensive training in the simulator. Yet managers are expected to fly their organizations relying on management school — the equivalent of ground school — and perhaps some experience as junior crew.

To meet these challenges we need to develop "management flight simulators," learning environments that motivate, that provide experiential as well as cognitive lessons, that compress time and space so that we may experience the long-term consequences of our actions. The Beer Game is one of a number of management flight simulators developed at MIT's Sloan School of Management for these purposes. The game was developed by Sloan's System Dynamics Group in the early 1960s as part of Jay Forrester's research on industrial dynamics. It has been played all over the world by thousands of people ranging from high school students to chief executive officers and government officials.

Of course, there is no beer in the Beer Game, and the game does not promote drinking. Originally called the "production-distribution game," the game was renamed because most students are more excited about producing beer than widgets or toasters. When played in, say, high schools, it easily becomes the apple juice game.

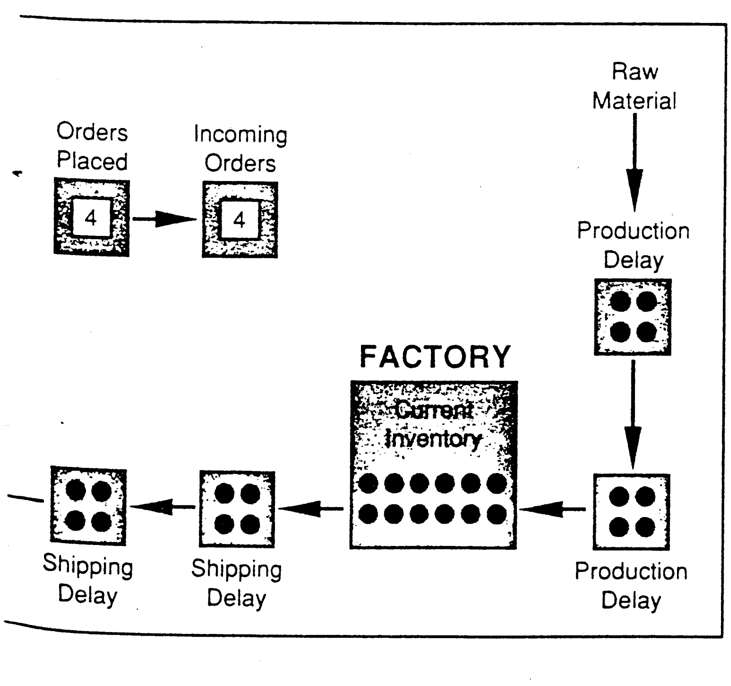

Playing the game

The game is played on a board that portrays the production and distribution of beer (Figures 1). Each team consists of four sectors: Retailer, Wholesaler, Distributor and Factory (R, W, D, F) arranged in a linear distribution chain. One or two people manage each sector. Pennies stand for cases of beer. A deck of cards represents customer demand. Each simulated week, customers purchase from the retailer, who ships the beer requested out of inventory. The retailer in turn orders from the wholesaler, who ships the beer requested out of their own inventory. Likewise, the wholesaler orders and receives beer from the distributor, who in turn orders and receives beer from the factory, where the beer is brewed. At each stage there are shipping delays and order processing delays. The players' objective is to minimize total team costs. Inventory holding costs are \$.50/case/week. Backlog costs are \$1.00/case/week, to capture both the lost revenue and the ill will a stockout causes among customers. Costs are assessed at each link of the distribution chain.

The game can be played with anywhere from four to hundreds of people. Each person is asked to bet \$1, with the pot going to the *team* with the lowest total costs, winner take all. The game is initialized in equilibrium. Each inventory contains 12 cases and initial throughput is four cases per week. In the first few weeks of the game the players learn the mechanics of filling orders, recording inventory, etc. During this time customer demand remains constant at four cases per week, and each player is directed to order four cases, maintaining the equilibrium. Beginning with week four the players are allowed to order any quantity they wish, and are told that customer demand may vary; one of their jobs is to forecast demand. Players are told the game will run for 50 simulated weeks, but play is actually halted after 36 weeks to avoid horizon effects.

Each player has good local information but severely limited global information. Players keep records of their inventory, backlog and orders placed with their supplier each week. However, people are directed not to communicate with one another; information is passed through orders and shipments. Customer demand is not known to any of the players in advance. Only the retailers discover customer demand as the game proceeds. The others learn only what their own customer orders.

These information limitations imply that the players are unable to coordinate their decisions or jointly plan strategy, even though the objective of each team is to minimize total costs. As in many real life settings, the global optimization problem must be factored into subproblems distributed throughout the organization.

The game is deceptively simple compared to real life. All you have to do is meet customer demand and order enough from your own supplier to keep your inventory low while avoiding costly backlogs. There are no machine breakdowns or other random events, no labor problems, no capacity limits or financial constraints. Yet the results are shocking.

Typical results: boom and bust

Figure 2 shows actual results from teams consisting of graduate students and business executives. Each column shows the results of a single team. The top four graphs show the orders placed by the players, from the retailer (bottom) to factory (top). The bottom four graphs show the players' inventories and backlogs (negative values), in the same order. Average team costs are about \$2,000, though it is not uncommon for costs to exceed \$10,000; few ever go below \$1,000. Optimal performance (calculated using only the information actually available to players themselves) is about

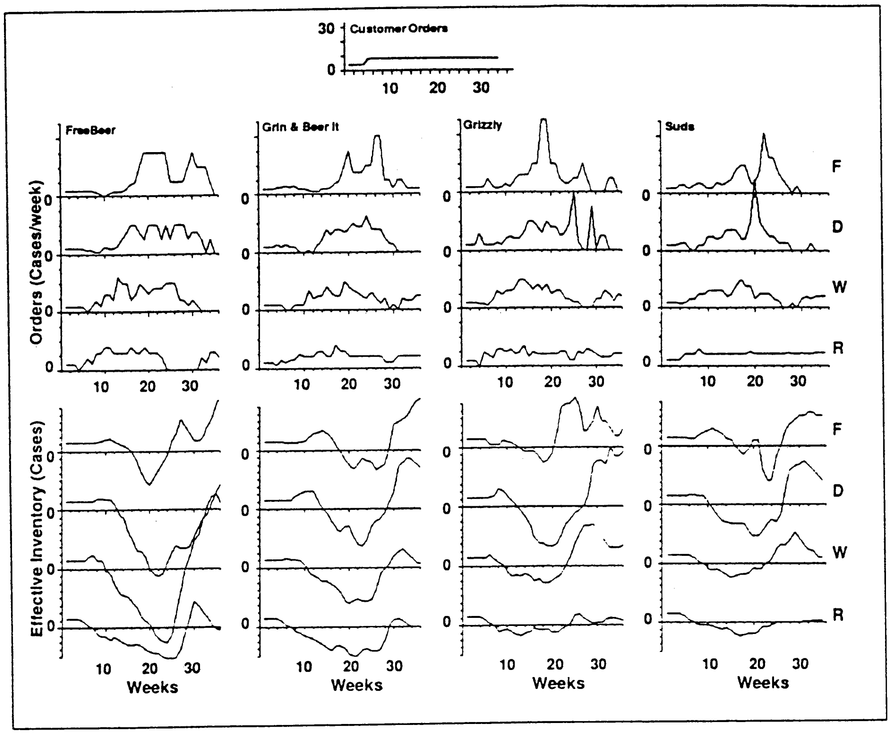

Figure 2: Typical Beer Game results. Top: orders; bottom: inventory (negative values denote backlogs). From bottom to top: Retailer, Wholesaler, Distributor, Factory. Tick marks denote 10 cases of beer. Compare the oscillations to the small step in customer orders.

$200. Average costs are *ten times* greater than optimal!

More revealing, the departures from optimality are not random. Though individual games differ quantitatively, they always exhibit the same patterns of behavior:

1. *Oscillation.* Orders and inventories are dominated by large amplitude fluctuations, with an average period of about 20 weeks.
2. *Amplification.* The amplitude and variance of orders increases steadily from customer to retailer to factory. The peak order rate at the factory is on average more than double the peak order rate at retail.
3. *Phase lag.* The order rate tends to peak later as one moves from the retailer to the factory.

In virtually all cases, the inventory levels of the retailer decline, followed in sequence by a decline in the inventory of the wholesaler, distributor and factory. As inventory falls, players tend to increase their orders. Players soon stock out. Backlogs of unfilled orders grow. Faced with rising orders and large backlogs, players dramatically boost the orders they place with their supplier. Eventually, the factory brews and ships this huge quantity of beer, and inventory levels surge. In many cases one can observe a second cycle.

Lessons of the game

During the game emotions run high. Many players report feelings of frustration and helplessness. Many blame their teammates for their problems; occasionally, heated arguments break out. After the game I ask the players to sketch their best estimate of the pattern of customer demand, that is, the contents of the customer order deck. Only the retailers have direct knowledge of that demand. The vast majority invariably draw a fluctuating pattern for customer demand, rising from the initial rate of four to a peak around 20 cases per week, then plunging.

"After all, it isn't my fault," people tell me, "if a huge surge in demand wiped out my stock and forced me to run a backlog. Then you tricked me – just when the tap began to flow, you made the customers go on the wagon, so I got stuck with all this excess inventory." Blaming the customer for the cycle is plausible. It is psychologically safe. And it is dead wrong. In fact, customer demand begins at four cases per week, then rises to eight cases per week in week five and remains completely constant ever after.

This revelation is often greeted by disbelief. How could the wild oscillations arise when the environment is virtually constant? Since the cycle isn't a consequence of fickle customers, players realize their own actions must have created the cycle. Though each player was free to make their own decisions, the same patterns of behavior emerge in every game, vividly demonstrating the powerful role of the system in shaping our behavior.

Research reported in Sterman [1989] shows how this occurs. Most people do not account for the impact of their own decisions on their teammates – on the system as a whole. In particular, people have great difficulty appreciating the multiple feedback loops, time delays and non-linearities in the system, using instead a very simple heuristic to place orders. When customer orders increase unexpectedly, retail inventories fall, since the shipment delays mean deliveries continue for several weeks at the old, lower rate. Faced with a growing backlog, people must order more than demand, often trying to fix the problem quickly by placing huge orders. If there were no time delays, this strategy would work well. But in the game, these large orders stock out the wholesaler. Retailers don't receive the beer they ordered, and grow increasingly anxious as their backlog worsens, leading them to order still more, even though the supply pipeline contains more than enough. Thus, the small step in demand from four to eight is amplified and distorted as it is passed to the wholesaler, who, reacting in kind, further amplifies the signal as it goes up the chain to the factory. Eventually, of course, the beer is brewed. The players cut orders as inventory builds up, but too late — the beer in the supply line continues to arrive. Inventories always overshoot, peaking at an average of about 40 cases.

Faced with what William James called the "bloomin', buzzin' confusion" of events, most people forget they are part of a larger whole. Under pressure, we focus on managing our own piece of the system, trying to keep our own costs low. And when the long-term effects of our short-sighted actions hit home, we blame our customer for ordering erratically, and our supplier for delivering late. Understanding how well-intentioned, intelligent people can create an outcome no one expected and no one wants is one of the profound lessons of the game. It is a lesson no lecture can convey.

The patterns of behavior observed in the game — oscillation, amplification and phase lag — are readily apparent in the real economy (Figure 3), from the business cycle to the recent boom and bust in real estate. The persistence of these cycles over decades is a major challenge to educators seeking to teach principles and tools for effective management. Though repeated experience with cycles in the real world should lead to learning and improvement, the duration of the business cycle exceeds the tenure of many managers. In real life the feedback needed to learn is delayed and confounded by change in dozens of other variables. By compressing time and space, and permitting controlled experimentation, management flight simulators can help overcome these impediments to learning from experience.

But the biggest impediments to learning are the mental models through which we construct our understanding of reality. By blaming outside forces we deny ourselves the opportunity to learn — recall that nearly all players conclude their roller coaster ride was caused by fluctuating demand. Focusing on external events leads people to seek better forecasts rather than redesigning the system to be robust in the face of the inevitable forecast errors. The mental models people bring to the understanding of complex dynamics systematically lead them away from the high leverage point in the system, hindering learning, and reinforcing the belief that we are helpless cogs in an overwhelmingly complex machine.

Thus to be effective, management flight simulators must be more than just business games. They must be embedded in a learning environment that encourages reflection on the perceptions, attributions and other mental models we use to interpret experience as well as the substantive lessons of the situation. These issues are the focus of current research at MIT and elsewhere [see Ster-

Under pressure, we focus on managing our own piece of the system, trying to keep our own costs low. And when the long-term effects of our short-sighted actions hit home, we blame our customer for ordering erratically, and our supplier for delivering late.

man and Morecroft 1992 for examples]. In addition to growing use in education, management flight simulators, both computer-based and manual, are finding successful application in a wide range of firms. They have helped stimulate improvement in hospital emergency room operations, raised maintenance productivity in the chemical industry, boosted service quality in the insurance industry, and helped top management in high-tech, petrochemicals and other industries to reformulate their strategies. Though much further work lies ahead, flight simulators may someday be as integral a part of the learning process for managers as they are today for pilots.

Using the Beer Game and other management flight simulators

The Beer Game is particularly useful in classes on operations management, production scheduling and related issues. The game highlights the importance of coordination among levels in an organization, the role of information systems in controlling complex systems, and the implications of different production paradigms such as just-in-time inventory management.

But the game illustrates more general lessons as well. The game creates a real organization, with "teams" supposed to work together. Yet the pressures of events and limited mental models of the players quickly cause team cohesion to break down. The game provides a vivid experience with a complex system, where players can see how the collective results of individually sensible decisions can be disastrous; where they can see the connection between the structure of a system and the dynamics it generates. The game is often used by firms in the service, financial and other industries where there is no inventory to manage. It is widely played as a team building experience at all levels of management from the shop floor to the boardroom.

Resources

A full analysis of the Beer Game appears in "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment" [Sterman, 1989]. Other management flight simulators and applications to real organizations are described in "Modelling for Learning" [Morecroft, Sterman, eds., 1992].

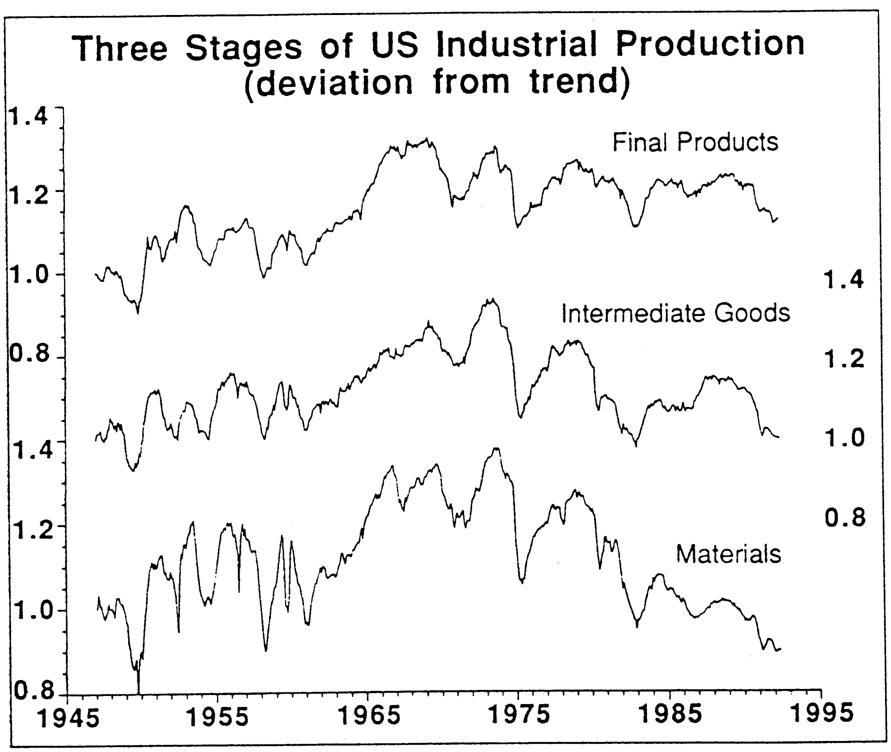

Figure 3. The 'beer game' in real life. The U.S. economy exhibits oscillation, amplification and phase lag as one moves through the distribution chain from production of consumer goods to intermediate goods to raw materials.

Instructions and a video tape of the Beer Game shown on the MacNeil-Lehrer News Hour in 1989 are available from the System Dynamics Group at MIT. A number of additional management flight simulators around other operational and strategic issues have also been developed at MIT. These simulators are computer-based and come with full documentation and instructions. The list includes:

- *People Express Airlines.* This computer simulation puts you in command of the innovative but now defunct People Express Airlines. You decide what prices to set, how fast to grow, how to respond to the competition. Your hiring policies influence morale, productivity and turnover; your marketing efforts shape the growth of demand; your competitors fight back. Widely used in marketing, strategy, organizational behavior, operations and even law schools.
- *B & B Enterprises.* You are responsible for the management of a new consumer durable product from launch through maturity. You set price, marketing budgets and build capacity as the product goes through its lifecycle. You must forecast demand for the product and respond to a simulated competitor in a dynamic world including learning curves, word of mouth, product differentiation, capacity acquisition lags and price conscious customers. This is useful in marketing, strategy, industrial organization, game theory, and modeling and simulation.

These simulations are currently available only for Macintosh computers. Contact John Sterman for information at the Sloan School of Management, Massachusetts Institute of Technology, 50 Memorial Drive, Cambridge, MA 02139. E-mail: jsterman@mit.edu.

References

1. Sterman, J. (1989), "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment," *Management Science*, Vol. 35, No. 3, pp. 321-339.
2. Morecroft, J., and Sterman, J., eds. (1992), "Modeling for Learning," special issue of the *European Journal of Operational Research*, Vol. 59 (May, 1992).

John Sterman is a professor at the Sloan School of Management at MIT.

MANAGEMENT SCIENCE
Vol. 35, No. 3, March 1989
Printed in U.S.A.

MODELING MANAGERIAL BEHAVIOR: MISPERCEPTIONS OF FEEDBACK IN A DYNAMIC DECISION MAKING EXPERIMENT*

JOHN D. STERMAN
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Studies in the psychology of individual choice have identified numerous cognitive and other bounds on human rationality, often producing systematic errors and biases. Yet for the most part models of aggregate phenomena in management science and economics are not consistent with such micro-empirical knowledge of individual decision-making. One explanation has been the difficulty of extending the experimental methods used to study individual decisions to aggregate, dynamic settings. This paper reports an experiment on the generation of macrodynamics from microstructure in a common managerial context. Subjects manage a simulated inventory distribution system which contains multiple actors, feedbacks, nonlinearities, and time delays. The interaction of individual decisions with the structure of the simulated firm produces aggregate dynamics which systematically diverge from optimal behavior. An anchoring and adjustment heuristic for stock management is proposed as a model of the subjects' decision processes. Econometric tests show the rule explains the subjects' behavior well. The estimation results identify several 'misperceptions of feedback' which account for the poor performance of the subjects. In particular, subjects are shown to be insensitive to the feedbacks from their decisions to the environment. Finally, the generality of the results is considered and implications for behavioral theories of aggregate social and economic dynamics are explored.
(BEHAVIORAL DECISION THEORY; DYNAMIC DECISION-MAKING; EXPERIMENTAL ECONOMICS; INVENTORY MANAGEMENT; SYSTEM DYNAMICS)

1. Introduction

Experimental studies in economics and the psychology of individual choice have identified numerous cognitive, informational, temporal, and other limitations which bound human rationality, often producing behavior which differs from the predictions of rational models (Simon 1979, Kahneman, Slovic, and Tversky 1982, Plott 1986, Smith 1986, Hogarth and Reder 1987). Yet for the most part models of aggregate phenomena in management science and economics are not consistent with such micro-empirical knowledge of individual decision-making. In a 1981 review Hogarth laments the "insufficient attention" paid "to the effects of feedback between organism and environment." By feedback is meant not merely outcome feedback but changes in the environment, in the conditions of choice, which are caused, directly and indirectly, by an agent's past actions. For example, a firm's decision to increase production feeds back through the market to influence the price of goods, profits, and demand; greater output may tighten the markets for labor and materials; competitors may react—all influencing future production decisions. Such multiple feedbacks are the norm rather than the exception in real problems of choice. Consequently, the focus of much research in behavioral decision theory on individual choice in static and discrete tasks has limited the penetration of psychological perspectives in theories of aggregate dynamics such as the behavior of firms, industries, and the economy. In response, many call for renewed empirical investigation designed to "secure new kinds of data at the micro level, data that will provide direct evidence about the behavior of economic agents and the ways in which they go about making their decisions" (Simon 1984, p. 40). Though crucial, securing such micro-level data is

* Accepted by Robert L. Winkler, former Departmental Editor; received September 21, 1987. This paper has been with the author 3 months for 1 revision.

0025-1909/89/3503/0321$01.25

not sufficient. Coleman (1987) argues that the greatest progress in coupling economics and psychology lies in understanding the "apparatus for moving from the level of the individual actor to the behavior of the system," that is, the generation of macrobehavior from microstructure.

This paper applies the experimental methods used so effectively in the study of individual behavior to the generation of macrodynamics from microstructure in a common managerial context. In the experiment subjects manage a simulated industrial production and distribution system, the "Beer Distribution Game". The decision-making task is straightforward: subjects seek to minimize total costs by managing their inventories appropriately in the face of uncertain demand. But the simulated environment is rich, containing multiple actors, feedbacks, nonlinearities, and time delays. The interaction of individual decisions with the structure of the simulated firm produces aggregate dynamics which diverge significantly and systematically from optimal behavior. An anchoring and adjustment heuristic for stock management is proposed as a model of the subjects' decision processes. Econometric tests show the rule explains the subjects' behavior well. Analysis of the results shows that the subjects fall victim to several 'misperceptions of feedback.' Specifically, subjects failed to account for control actions which had been initiated but not yet had their effect. Subjects were insensitive to feedbacks from their decisions to the environment. The majority attributed the dynamics they experienced to external events, when in fact these dynamics were internally generated by their own actions. Further, the subjects' open-loop mental model, in which dynamics arise from exogenous events, is hypothesized to hinder learning and retard evolution towards greater efficiency. Finally, the generality of the results is considered and implications for behavioral theories of aggregate social and economic dynamics are discussed.

2. The Stock Management Problem

One of the most common dynamic decision-making tasks is the regulation of a stock or system state. In such a task, the manager seeks to maintain a quantity at a particular target level, or at least within an acceptable range. Stocks cannot be controlled directly but rather must be influenced by changes in their inflow and outflow rates. Typically, the manager must set the inflow rate so as to compensate for losses and usage and to counteract disturbances which push the stock away from its desired value. Often there are lags between the initiation of a control action and its effect, and/or lags between a change in the stock and the perception of that change by the decision maker. The duration of these lags may vary and may be influenced by the manager's own actions.

Stock management problems occur at many levels of aggregation. At the level of a firm, managers must order parts and raw materials so as to maintain inventories sufficient for production to proceed at the desired rate, yet prevent costly inventories from accumulating. They must adjust for variations in the usage and wastage of these materials and for changes in their delivery delays. At the level of the individual, people regulate the temperature of the water in their morning shower, guide their cars down the highway, and manage their checking account balances. At the macroeconomic level, the Federal Reserve seeks to manage the stock of money to stimulate economic growth and avoid inflation, while compensating for variations in credit demand, budget deficits, and international capital flows.

The generic stock management control problem may be divided into two parts: (i) the stock and flow structure of the system; and (ii) the decision rule used by the manager (Figure 1). Considering first the stock and flow structure, the stock S is the accumulation of the acquisition rate A less the loss rate L:

$$S_t = \int_{t_0}^{t} (A_\tau - L_\tau)d\tau + S_{t_0}. \tag{1}$$

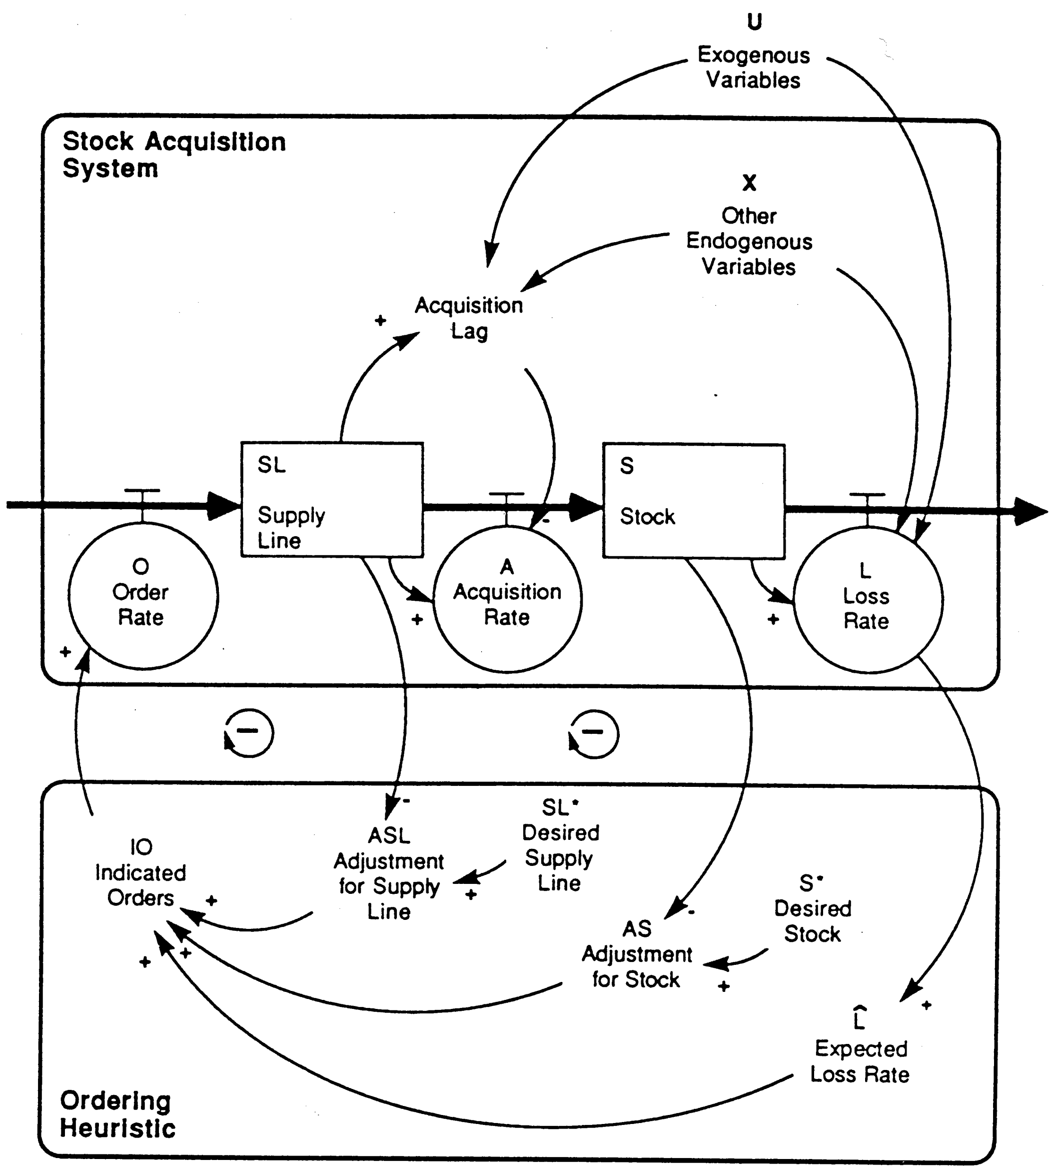

FIGURE 1. The Generic Stock-Management System.
Rectangles denote state variables; heavy arrows and 'valves' denote rates of flow (see equations (1)–(2)). The polarity of the information feedbacks denotes the sign of the relationship between independent and dependent variables, e.g., $X \rightarrow {}^{+}Y \Rightarrow (\partial Y/\partial X) > 0$.

Losses here include any outflow from the stock and may arise from usage (as in a raw material inventory) or decay (as in the depreciation of plant and equipment). The loss rate must depend on the stock itself—losses must approach zero as the stock is depleted—and may also depend on other endogenous variables **X** and exogenous variables **U**. Losses may be nonlinear and may depend on the age distribution of the stock.

The acquisition rate depends on the supply line *SL* of units which have been ordered but not yet received, and the average acquisition lag λ. In general, λ may depend on the supply line itself and on the other endogenous and exogenous variables. The supply line is simply the accumulation of the orders which have been placed *O* less those which have been delivered:

$$SL_t = \int_{t_0}^{t} (O_\tau - A_\tau)d\tau + SL_{t_0}. \tag{2}$$

The structure represented by Figure 1 and equations (1)–(2) is quite general. The system may be nonlinear. There may be arbitrarily complex feedbacks among the endogenous variables, and the system may be influenced by a number of exogenous forces, both systematic and stochastic. Table 1 maps common examples into the generic form. In each case, the manager must choose the order rate over time so as to keep the stock close to a target. It is interesting to note that the characteristic behavior modes of many of these systems include oscillation and instability.

In most realistic stock management situations the complexity of the feedbacks among the variables precludes the determination of the optimal strategy. The order decision model proposed here assumes that managers, unable to optimize, instead exercise control through a heuristic which is locally rational. The model thus falls firmly in the tradition of bounded rationality as developed by Simon (1982), Cyert and March (1963), and others. Cognitive limitations are recognized, as are information limitations caused by organizational structures such as task factoring and subgoals (for a discussion of local rationality in the context of simulation models see Morecroft 1983, 1985 and Sterman 1985, 1987a).

The hypothesized decision rule utilizes information locally available to the decision maker and does not presume that the manager has global knowledge of the structure of the system. Managers are assumed to choose orders so as to: (1) replace expected losses from the stock; (2) reduce the discrepancy between the desired and actual stock; and (3) maintain an adequate supply line of unfilled orders. To formalize this heuristic, first observe that orders in most real-life situations must be nonnegative:

$$O_t = \text{MAX}\,(0, IO_t) \tag{3}$$

where IO is the indicated order rate, the rate indicated by other pressures. Order cancellations are sometimes possible and may sometimes exceed new orders (e.g. the U.S. nuclear power industry in the 1970s). Cancellations are likely to be subject to different costs and administrative procedures than new orders and should be modeled as a distinct outflow from the supply line rather than as negative orders.

The indicated order rate is based on the anchoring and adjustment heuristic (Tversky and Kahneman 1974). Anchoring and adjustment is a common strategy in which an unknown quantity is estimated by first recalling a known reference point (the anchor) and then adjusting for the effects of other factors which may be less salient or whose effects are obscure, requiring the subject to estimate these effects by what Kahneman and Tversky (1982) call 'mental simulation.' Anchoring and adjustment has been shown to apply to a wide variety of decision-making tasks (Einhorn and Hogarth 1985, Davis et al. 1986, Johnson and Schkade 1987, Hines 1987). Here the anchor is the expected loss rate $\hat{L}$. Adjustments are then made to correct discrepancies between the desired and actual stock (AS), and between the desired and actual supply line (ASL):

$$IO_t = \hat{L}_t + AS_t + ASL_t. \tag{4}$$

Expected losses may be formed in various ways. Common formulations include static expectations $\hat{L}_t = L^*$ (a constant or equilibrium value), regressive expectations $\hat{L}_t = \gamma L_{t-1} + (1 - \gamma)L^*$, $0 \le \gamma \le 1$, adaptive expectations $\hat{L}_t = \theta L_{t-1} + (1 - \theta)\hat{L}_{t-1}$, $0 \le \theta \le 1$, and extrapolative expectations, $\Delta\hat{L}_t = \sum \omega_i \cdot \Delta L_{t-i}$, where Δ is the first difference operator and $\omega_i \ge 0$.

The feedback structure of the heuristic is shown in the bottom part of Figure 1. The adjustment for the stock AS creates a negative feedback loop which regulates the stock. For simplicity the adjustment is linear in the discrepancy between the desired stock S^* and the actual stock:

$$AS_t = \alpha_S(S_t^* - S_t), \tag{5}$$

TABLE 1
Examples of Stock-Management Systems

System	Stock	Supply Line	Loss Rate	Acquisition Rate	Order Rate	Typical Behavior
Inventory Management	Inventory	Goods on Order	Shipments to Customers	Arrivals from supplier	Orders for goods	Business cycles
Capital investment	Capital Plant	Plant under construction	Depreciation	Construction completion	New contracts	Construction cycles
Equipment	Equipment	Equipment on order	Depreciation	Equipment delivery	New equipment orders	Business cycles
Human Resources	Employees	Vacanies & trainees	Layoffs and quits	Hiring rate	Vacancy creation	Business cycles
Cash Management	Cash balance	Pending loan applications	Expenditures	Borrowing rate	Loan application rate	?
Marketing	Customer Base	Prospective customers	Defections to competitors	Recruitment of new customers	New customer contacts	?
Hog farming	Hog stock	Immature and gestating hogs	Slaughter rate	Maturation rate	Breeding rate	Hog cycles
Agricultural commodities	Inventory	Crops in the field	Consumption	Harvest rate	Planting rate	Commodity cycles
Commercial construction	Building stock	Buildings under development	Depreciation	Completion rate	Development rate	15–25 year cycles
Cooking on electric range	Temperature of pot	Heat in coils of range	Diffusion to air	Diffusion from coils to pot	Setting of burner	Overcooked dinner
Driving	Distance to next car	Momentum of car	Friction	Velocity	Gas and Brake pedals	Stop-and-go traffic
Showering	Water Temperature	Water Temp. in pipes	Drain rate	Flow from showerhead	Faucet settings	Burn-then-freeze
Personal energy level	Glucose in bloodstream	Sugar and starch in GI tract	Metabolism	Digestion	Food consumption	Cycles of energy level
Social drinking	Alcohol in blood	Alcohol in stomach	Metabolism of alcohol	Diffusion from stomach to blood	Drinking rate	Drunkenness

where the stock adjustment parameter α_S is the fraction of the discrepancy ordered each period. The adjustment for the supply line is formulated analogously as

$$ASL_t = \alpha_{SL}(SL_t^* - SL_t), \tag{6}$$

where SL^* is the desired supply line and α_{SL} is the fractional adjustment rate for the supply line. The desired supply line in general is not constant but depends on the desired throughput Φ^* and the expected lag between ordering and acquisition of goods:

$$SL_t^* = \hat{\lambda}_t \cdot \Phi_t^*. \tag{7}$$

The longer the expected delay in acquiring goods or the larger the throughput desired, the larger the supply line must be. For example, if a retailer wishes to receive 1,000 widgets per week from the supplier and delivery requires 6 weeks, the retailer must have 6000 widgets on order to ensure an uninterrupted flow of deliveries. The adjustment for the supply line creates a negative feedback loop which adjusts orders so as to maintain an acquisition rate consistent with the desired throughput and the acquisition lag. Without such a feedback orders would be placed even after the supply line contained sufficient orders to correct stock shortfalls, producing overshoot and instability. The supply line adjustment also compensates for changes in the acquisition lag. If the acquisition lag doubled, for example, the supply line adjustment would induce sufficient additional orders to restore the desired throughput. As in the formation of expected losses, there are a variety of possible representations for $\hat{\lambda}$ and Φ^*, ranging from constants through sophisticated forecasts.

In terms of anchoring and adjustment, expected losses form an easily anticipated and relatively stable starting point for the determination of orders. Loss rate information will typically be locally available and highly salient to the decision maker. Replacing losses will keep the stock constant at its current level. Adjustments are then made in response to the adequacy of the stock and supply line. No assumption is made that these adjustments are optimal or that managers actually calculate the order rate using the equations (Einhorn, Kleinmuntz, and Kleinmuntz 1979). Rather, pressures arising from the discrepancies between desired and actual quantities cause managers to adjust the order rate above or below the level that would maintain the status quo.

3. A Stock Management Experiment

The "Beer Distribution Game" is a role-playing simulation of an industrial production and distribution system developed at MIT to introduce students of management to the concepts of economic dynamics and computer simulation. In use for nearly three decades, the game has been played all over the world by thousands of people ranging from high school students to chief executive officers and government officials.

The game is played on a board which portrays the production and distribution of beer (Figure 2). Orders for and cases of beer are represented by markers and pennies which are manipulated by the players. Each brewery consists of four sectors: retailer, wholesaler, distributor, and factory (R, W, D, F). One person manages each sector. A deck of cards represents customer demand. Each week, customers demand beer from the retailer, who ships the beer requested out of inventory. The retailer in turn orders beer from the wholesaler, who ships the beer requested out of the wholesaler's inventory. Likewise the wholesaler orders and receives beer from the distributor, who in turn orders and receives beer from the factory. The factory produces the beer. At each stage there are shipping delays and order receiving delays. These represent the time required to receive, process, ship, and deliver orders, and as will be seen play a crucial role in the dynamics.

The subjects' objective is to minimize total company costs during the game. Inventory holding costs are \$.50/case/week, and stockout costs (costs for having a backlog of

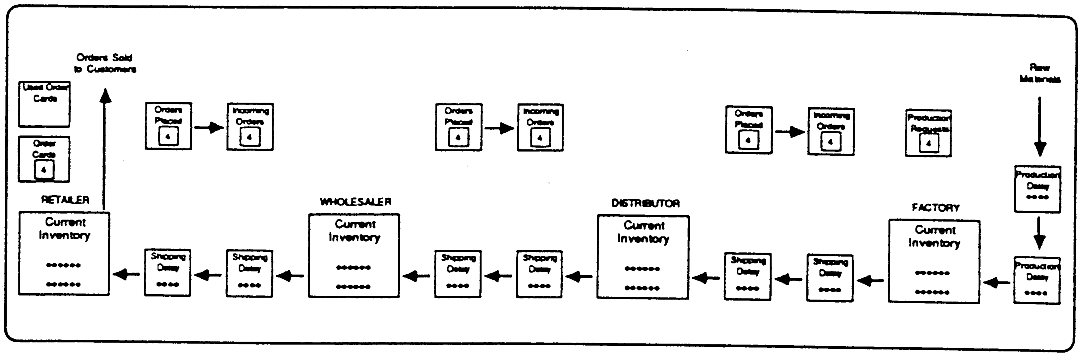

FIGURE 2. "Beer Distribution Game" Board.

Initial conditions are shown: each inventory contains 12 pennies; each shipping/production delay contains 4. Orders are 4 throughout the distribution chain. During actual play the order cards are face down at all times. Each simulated week requires all subjects to carry out five steps:

1. *Receive inventory and advance shipping delays*. The contents of the shipping delay immediately to the right of the inventory are added to the inventory; the contents of the shipping delay on the far right are moved into the delay on the near right. The factory advances the production delays.
2. *Fill orders*. Retailers take the top card in the customer order deck, others examine the contents of "Incoming Orders". Orders are always filled to the extent inventory permits. Unfilled orders add to the backlog, if any. The number of orders to fill is the incoming order plus any backlog from the prior week.
3. *Record inventory or backlog on the record sheet*.
4. *Advance the order slips*. Order slips in the "Orders Placed" box are moved to the "Incoming Orders" box on the immediate right. Factories introduce the contents of "Production Requests" into the top production delay.
5. *Place orders*. Each player decides what to order, records the order on the record sheet and on an order slip which is placed face down in the "Orders Placed" box. Factories place their orders in "Production Requests."

Note that only step 5, Place Orders, involves a decision on the part of the subject. Steps 1–4 handle bookkeeping and other routine tasks.

unfilled orders) are \$1.00/case/week. Costs are assessed at each link of the distribution chain.

The decision task of each subject is a clear example of the stock management problem. Subjects must keep their inventory as low as possible while avoiding backlogs. Inventory must be ordered, and the delivery lag is potentially variable (that lag is never less than 4 weeks but may be longer if upstream inventories are insufficient).

Experimental Protocol

Typical sessions involve three to eight teams of four players. Subjects are randomly assigned roles as retailer, wholesaler, etc. Each subject is asked to place \$1 in a kitty to be wagered against the other teams. The kitty goes to the team with the lowest total costs, winner take all.[1] Next, the steps of the game are explained (Figure 2). The game is initialized in equilibrium. Each inventory contains 12 cases and initial throughput is four cases per week (Figure 2). Customer demand likewise begins at four cases per week. The first four weeks of play are used to familiarize the subjects with the mechanics of filling orders, recording inventory, etc. During this time customer demand remains constant, and each player is directed to order four cases, maintaining the initial equilibrium. Beginning with week four the players are allowed to order any nonnegative quantity they wish. There is an unannounced, one-time increase in customer demand to eight cases

[1] Protocols for experimental economics (e.g. Smith 1982) call for monetary rewards geared to performance. However, a number of experiments have shown performance is not significantly improved and may be worsened by higher reward levels (e.g. Grether and Plott 1979, Slovic and Lichtenstein 1983, Tversky and Kahneman 1981). Here subjects wager \$1 for a chance to win about \$4. Though small, these rewards emphasize the goal of minimum *team* costs and appear to have a powerful motivating effect.

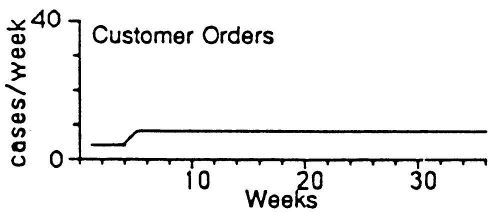

FIGURE 3. Customer Orders.
Customer orders rise from 4 to 8 cases per week in week 5. Vertical tick-marks denote 10 units. Compare against the subjects' orders (Figure 4).

per week in week 5 (Figure 3). The step creates a disequilibrium disturbance to which the subjects must react while facilitating subsequent analysis.

During the sessions questions concerning rules, procedures, or interpretation are answered; questions concerning strategy or customer demand are not. Subjects are told the game will run for 50 simulated weeks, but play is actually halted after 36 weeks to avoid horizon effects. Typically the game is introduced and played in 90 minutes, followed by a debriefing session.

Information Availability

The game is designed so that each subject has good local information but severely limited global information. Each maintains a record sheet which includes their inventory or backlog and orders placed with their supplier for each week. However, subjects are directed not to communicate with one another, either across or within a game. Customer demand is not known to any of the subjects in advance. Retailers are the only subjects who discover customer demand as the game proceeds. The others learn only what their own customer orders, and only after a delay of one week. The players do sit next to one another, and some crosstalk is unavoidable. Each can readily inspect the board to see how large the inventories of beer are at the other stations, thus gleaning information potentially useful in ordering. Game play is usually quite lively and the subjects' outbursts may also convey information.

These information limitations imply that the subjects are unable to coordinate their decisions or jointly plan strategy, even though the objective of each team is to minimize total costs. As in many real situations, the problem of global optimization must be factored into subgoals which are distributed throughout the organization.

The Sample

The results reported here were drawn from 48 trials (192 subjects) collected over a period of four years. Since the subjects keep the records manually there are occasional accounting errors. Trials in which any of the four subjects made significant errors were discarded. Eleven trials were retained (44 subjects). That sample consists of undergraduate, MBA, and Ph.D. students at MIT's Sloan School of Management, executives from a variety of firms participating in short courses on computer simulation, and senior executives of a major computer firm. Analysis showed the trials with the highest costs to be most prone to accounting errors. Thus the final sample of eleven is biased towards those who understood and performed best in the game. The effect is modest, however, and reinforces the conclusions drawn below.

4. Results

The complexity of the system—it is a 23rd order nonlinear difference equation—renders calculation of the optimal behavior intractable. However, a benchmark for evaluating the performance of the subjects was obtained through computer simulation. As

TABLE 2

Comparison of Experimental and Benchmark Costs. Benchmark costs are the minimum costs produced by simulation of the proposed decision rule and are an upper bound estimate of optimal performance in the experiment

	Team Total	Retailer	Wholesaler	Distributor	Factory
Mean ($N = 11$)	\$2028	\$383	\$635	\$630	\$380
Benchmark	\$204	\$46	\$50	\$54	\$54
Ratio	9.9	8.3	12.7	11.7	7
t-statistic:	8.7	4.9	5.9	6.9	9.7
H_0: Mean cost = Benchmark	$p < 0.000+$	$p < 0.001$	$p < 0.000+$	$p < 0.000+$	$p < 0.000+$

implemented below, the proposed decision rule involves four parameters. The parameters which produce minimum total costs were calculated by simulation of the game over the plausible parameter space.[2] The benchmark costs were computed subject to the same information limitations faced by the subjects. Benchmark costs are shown in Table 2 compared to actual costs for the eleven trials. The average team cost is ten times greater than the benchmark. The individual sectors exceed the benchmark costs by similar ratios. The differences between actual and benchmark costs are highly significant.

More interesting is the character of the departures from optimality. Are the subjects behaving in similar ways? Do their errors arise from common sources? Figure 4 shows several typical trials; Table 3 summarizes key indicators of the behavior for the full sample. Examination of the order pattern reveals several regularities.

1. *Oscillation*. The trials are all characterized by instability and oscillation. Orders and inventory are dominated by large amplitude fluctuations, with an average of 21 weeks required to recover initial inventory levels. In virtually all cases, the inventory levels of the retailer decline, followed in sequence by a decline in the inventory of the wholesaler, distributor, and factory (Figure 4). As inventory falls, subjects tend to increase their orders. 'Effective inventory' (inventory less any backlog of unfilled orders) generally becomes significantly negative, indicating the sectors have backlogs. The maximum backlog averages 35 cases, and occurs between weeks 20 and 25. As additional product

TABLE 3

Summary of Experimental Results. Averages of 11 Trials

	Customer	Retailer	Wholesaler	Distributor	Factory
Periodicity (weeks)					
Time to recover initial inventory	N/A	24	23	22	16
Date of Minimum Inventory	N/A	20	22	20	22
Date of Maximum Inventory	N/A	28	27	30	26
Amplification					
Peak Order Rate (cases/week)	8	15	19	27	32
Variance of Order Rate (cases/week)^2	1.6	13	23	45	72
Peak Inventory (cases)	N/A	20	41	49	50
Minimum Inventory (cases)	N/A	−25	−46	−45	−23
Range (cases)	N/A	45	88	94	73
Phase Lag					
Date of Peak Order Rate (week)	5	16	16	21	20

[2] To reduce the search space the same parameters are used in each sector. The optimal parameters are $\theta = 0$, $\alpha_S = 1$, $\beta = 1$, and $S' = 28$ (20 for the factory).

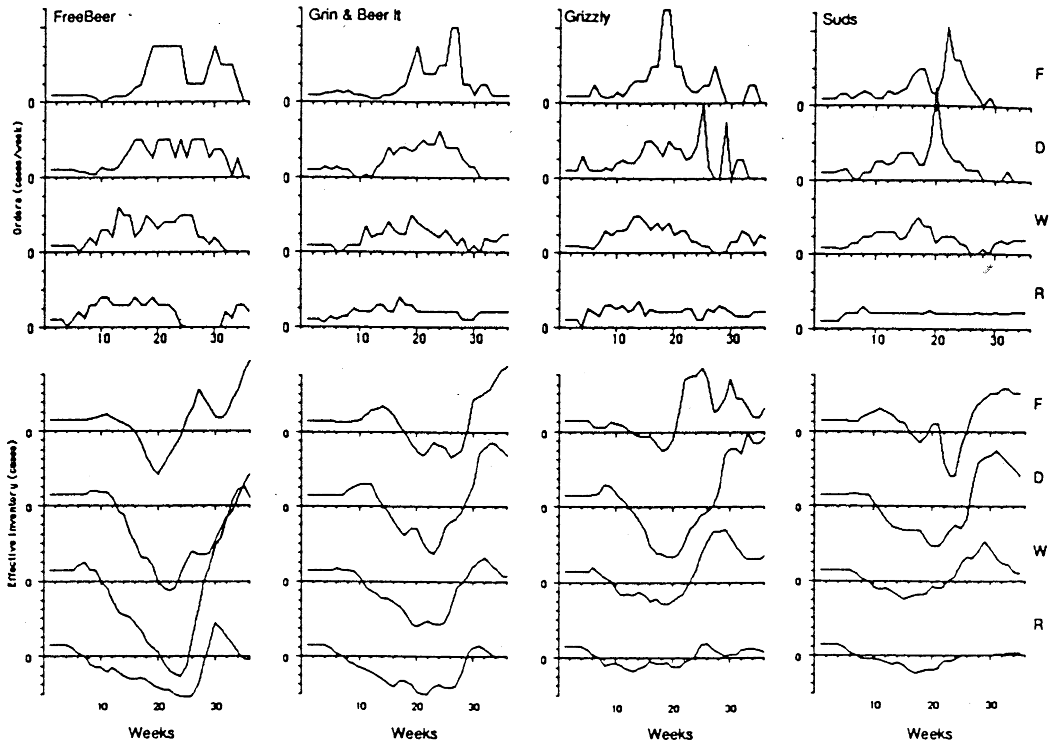

FIGURE 4. Experimental Results for Four Typical Trials.
Top: Orders; bottom: inventory (from bottom to top, Retailer; Wholesaler, Distributor, Factory). Tick-marks on y-axes denote 10 units. Note the oscillation, amplification, and phase lag as the change in customer orders propagates from retailer to factory.

is brewed and shipped inventory levels surge. Inventory in many cases substantially overshoots its initial levels. The inventory peak averages 40 cases and occurs between weeks 25 and 30. Orders fall off rapidly as inventory builds up.

2. *Amplification*. The amplitude and variance of orders increases steadily from customer to retailer to factory. The peak order rate at the factory is on average more than double the peak order rate at retail. Customer orders increase from 4 to 8 cases per week; by the time the disturbance has propagated to the factory the order rate *averages* a peak of 32 cases, an amplification factor of 700%.[3] Amplification in inventory excursions is also apparent. Note that the average period and excursion of factory inventories are somewhat less than those of the distributor and wholesaler. The factory, as primary producer, faces a shorter and constant delay in acquiring beer and can therefore correct inventory discrepancies faster and more reliably than the other sectors. This subtlety in the outcomes illustrates the extent to which the feedback structure of the task shapes the behavior of the subjects.

3. *Phase lag*. The order rate tends to peak later as one moves from the retailer to the factory. Customer orders increase from 4 to 8 in week 5. Retailer orders do not reach their peak until week 16, on average. Factory orders lag behind still further, peaking at week 20 on average. The phase lag is not surprising since the disturbance in customer orders must propagate through decision-making and order delays from retailer to wholesaler and so on.[4]

[3] Amplification is a rough measure of closed-loop gain and is measured as the excursion in the output variable relative to that of the input, in this case Δ(Factory Orders)/Δ(Customer Orders) = $(32 - 4)/(8 - 4) = 7$.

[4] There is no apparent lag between retailer and wholesaler or between distributor and factory, perhaps indicating that subjects used information outside their own sector.

Thus while the behavior of the subjects is plainly far from optimal, their behavior exhibits significant regularities, suggesting the subjects used similar heuristics to determine their orders. The pervasiveness and qualitative similarity of the oscillations is particularly noteworthy since the customer order rate, the only external disturbance, does not oscillate and is in fact virtually constant. The oscillation is endogenously produced by the interaction of the subjects' decisions with the feedback structure of the system. Explaining the origin of the cycle and the determinants of its period and amplitude are major tasks for any theory of dynamic decision-making behavior.

5. Testing the Theory

The decision rule must next be adapted to the particulars of the beer game and cast in a form suitable for estimation of the parameters. In the experiment, the stock S corresponds to the effective inventory of the subject and the supply line SL to the sum of orders in the mail delays, the backlog of the subject's supplier (if any), and the beer in the shipping delays. The loss rate is the rate at which each subject receives orders. To test the rule it is necessary to specify expected losses $\hat{L}$, the desired stock S^*, and the desired supply line SL^*.

Expected losses from the stock are the rate at which each subject expects their immediate customer to place orders, that is, the retailer's forecast of the customer order rate, the factory's forecast of the distributor's order rate, etc. Adaptive expectations are postulated. Adaptive expectations are widely used in simulation modeling of economic systems, are often a good model of the evolution of expectations in the aggregate (Sterman 1987b, Frankel and Froot 1987), and are one of the simplest formulations for expectations suitable for nonstationary processes.

Theory suggests the desired stock should be chosen to minimize expected costs given the cost function and expected variability of deliveries and incoming orders. However, the subjects have neither the time nor information to determine optimal inventory levels. The asymmetry of the cost function does suggest desired inventory should be greater than zero. In the absence of a procedure to calculate optimal inventory levels, however, one might expect the subjects' choice of S^* to be anchored to the initial level of 12 units. This hypothesis is tested below.

In general the desired supply line is variable and depends on the anticipated delay in receiving orders. However, subjects lack the means to determine the current lag in receiving orders. That lag is never less than four weeks but may be longer if the supplier has insufficient inventory to fill incoming orders. The desired supply line SL^* is therefore assumed to be constant.

The generic decision rule of equations (3)–(7) then becomes:

$$O_t = \text{MAX}\,(0, \hat{L}_t + AS_t + ASL_t), \tag{8}$$

$$\hat{L}_t = \theta L_{t-1} + (1-\theta)\hat{L}_{t-1}, \qquad 0 \le \theta \le 1, \tag{9}$$

$$AS_t = \alpha_S(S^* - S_t), \tag{10}$$

$$ASL_t = \alpha_{SL}(SL^* - SL_t), \tag{11}$$

where S^* and SL^* are constants. Defining $\beta = \alpha_{SL}/\alpha_S$ and $S' = S^* + \beta SL^*$, collecting terms, and allowing for an additive disturbance term ϵ yields

$$O_t = \text{MAX}\,[0, \hat{L}_t + \alpha_S(S' - S_t - \beta SL_t) + \epsilon_t]. \tag{12}$$

Note that since S^*, SL^*, α_{SL} and α_S are all ≥ 0, $S' \ge 0$. Further, subjects are unlikely to place more emphasis on the supply line than on inventory itself: the supply line does not directly enter the cost function nor is it as salient as inventory. Therefore it is probable that $\alpha_{SL} \leq \alpha_S$, meaning $0 \leq \beta \leq 1$. Thus β can be interpreted as the fraction of the supply line taken into account. If $\beta = 1$, the subjects fully recognize the supply line and do not overorder. If $\beta = 0$, goods on order are ignored.

The decision rule contains four parameters to be estimated (θ, α_S, S', and β) and is nonlinear. The disturbance ϵ is assumed to be Gaussian white noise. In this case, maximum likelihood estimates are found by minimizing the sum of squared errors $\Sigma\, e_t^2$. The estimated parameters of such nonlinear models are consistent and asymptotically efficient, and the usual measures of significance such as the t-test are asymptotically valid (Judge et al. 1980).[5] The Durbin-Watson test showed no significant residual autocorrelation for 23 of 44 subjects. Monte Carlo simulations showed the estimation procedure was not significantly degraded by autocorrelation in the disturbance as high as $\rho = 0.9$.

Table 4 shows the estimated parameters together with R^2 and root mean square errors. The mean R^2 is 71%; R^2 is less than 50% for only 6 of 44 subjects. A large majority of the estimated parameters are significant. Only 7 values of α_S, 4 values of S', and 15 values of β are not significantly different from zero. Of course any of these parameters could legitimately take on a value of zero. Zero is in fact the estimated value for 14 of the 26 insignificant estimates, and the standard errors of these estimates are smaller, on average, than those for the rest of the sample. However, two-thirds of the estimated values of θ are not significant. It appears that there is insufficient variation in incoming orders to determine if the expectation formation process is misspecified for these subjects.[6]

As a further test the game was simulated using the decision rule with the estimated parameters for each sector. Note that the costs incurred by a sector depend not only on the behavior of that sector but on all the other sectors in the distribution chain, and thus on the vectors of parameters $\boldsymbol{\theta}$, $\boldsymbol{\alpha_S}$, $\boldsymbol{S'}$, and $\boldsymbol{\beta}$ for the entire chain. If the rule were perfect, simulated and actual costs would be equal, and regression of the simulated costs on the actual costs would produce a slope of unity (t-statistic in parentheses):

$$\text{Costs}_{i,j} = \underset{(16.7)}{1.11} * \text{Simulated Costs}\,(\theta_j, \alpha_{S_j}, S'_j, \beta_j)_i; \quad i = R, W, D, F; \quad j = 1, \ldots, 11,$$

$N = 44, \qquad R^2 = 0.40.$

The slope is less than two standard errors from unity and highly significant, indicating an excellent correspondence between actual and simulated costs.

There is, however, a modest bootstrapping effect. Replacing the subjects with the model of their behavior improves performance. The average improvement is about 5% of actual costs. The improvement arises from the consistency of the decision rule compared to the subjects, who often changed orders from week to week, introducing high-frequency noise (Figure 4). The magnitude of the bootstrapping effect is comparable to that found in many prior studies of bootstrapping (reviewed in Camerer 1981) even though these studies involved linear models of clinical judgments where there were in general no significant

[5] Estimates were found by grid search of the parameter space subject to the constraints $0 \leq \theta \leq 1$ and α_S, S', $\beta \geq 0$. θ, α_S, β, and S' were estimated to the nearest 0.1, 0.05, 0.05, and 1 units, respectively. The search space was large enough to ensure capturing the global minimum of $\Sigma\, e_t^2$. The data and computer programs are available from the author. Because the ordering function does not contain a regression constant, the residuals need not satisfy $\Sigma\, e_t = 0$ (estimated and actual orders need not have a common mean) and the conventional R^2 is not an appropriate measure of fit. The alternative $R^2 = r^2$ is used, where r is the simple correlation between estimated and actual orders (Judge et al. 1980).

[6] θ can only be identified if L_t and $\hat{L}_t$ differ. Since $\hat{L}_t$ approaches L_t over time, a tight estimate of θ requires large variation in incoming orders from period to period. For all the retailers and several other sectors the variation in incoming orders is slight (recall that retailers face virtually constant demand). In fact, the 6 largest standard errors for θ are retailers. The hypothesis that expectations of customer demand adapt to past orders for these subjects cannot therefore be rejected; for one third of the sample it is supported.

TABLE 4

Estimated Parameters

Trial & Position	θ	α_s	β	S'	R^2	RMSE
Bassbeer						
R	0.90	0.10	0.65 a	20 a	0.20	3.13
W	0.00	0.25 a	0.50 a	27 a	0.86	1.99
D	0.15	0.05 a	0.35	14	0.74	2.76
F	1.00 a	0.65 a	0.40 a	15 a	0.84	4.56
Budweiser						
R	0.00	0.40 a	0.10 a	7 a	0.67	2.60
W	0.00	0.40 a	0.75 a	30 a	0.92	1.32
D	0.00	0.30 a	0.10 a	10 a	0.88	2.09
F	0.25 c	0.25 a	0.10	9 a	0.87	2.52
Coors						
R	0.00	0.20 a	0.00	25 a	0.57	1.60
W	0.00	0.15 a	0.50 a	38 a	0.11	2.84
D	0.90 a	0.30 a	0.20 a	10 a	0.61	2.84
F	0.25	0.30 a	0.00	18 a	0.73	4.07
Freebeer						
R	0.40	0.35 a	0.45 a	15 a	0.43	4.29
W	0.30	0.05 a	0.00	30 c	0.76	3.57
D	0.05	0.35 a	1.00 a	18 a	0.86	2.72
F	0.25	0.25 a	0.00	19 a	0.89	3.82
Grin & Beer It						
R	0.10	0.35 a	0.65 a	13 a	0.60	1.79
W	0.95 a	0.15 a	0.55 a	14 a	0.79	2.24
D	0.20 b	0.20 a	0.30 a	19 a	0.94	1.75
F	0.25	0.35 a	0.55 a	24 a	0.73	5.02
Grizzly						
R	0.05	0.30 a	0.65 a	31 a	0.58	1.88
W	0.30	0.20 a	0.35 a	27 a	0.82	2.32
D	0.15	0.05	0.25	15	0.32	7.47
F	0.55 a	0.65 a	0.00	9 a	0.75	5.93
Heineken1						
R	0.95	0.15 a	0.00	9 a	0.75	1.92
W	0.50 a	0.00	N/D	N/D	0.87	1.25
D	0.20 a	0.30 a	0.05 a	8 a	0.98	0.96
F	0.80 b	0.00	N/D	N/D	0.60	3.70
Heineken2						
R	0.50	0.05	0.60	6	0.10	4.08
W	0.40 a	0.10 a	0.30 a	16 a	0.81	2.18
D	1.00 a	0.15 a	0.80 a	14 a	0.73	3.26
F	0.55 a	0.80 a	0.00	9 a	0.87	3.08
Heineken3						
R	0.05	0.30 a	0.45 a	5 a	0.89	0.97
W	0.20	0.00	N/D	N/D	0.23	3.17
D	0.30 a	0.10 a	0.90 a	12 a	0.94	0.83
F	0.00	0.30 a	0.15 c	17 a	0.87	1.46
Suds						
R	1.00	0.00	N/D	N/D	0.76	0.85
W	0.05	0.30 a	0.20 a	20 a	0.76	2.23
D	0.15	0.60 a	0.35 a	0	0.69	5.19
F	0.40 a	0.35 a	1.05 a	32 a	0.95	2.06
Twoborg						
R	0.75	0.35 a	0.00	4 a	0.83	1.53
W	0.00	0.25 a	0.05	18 a	0.72	2.65
D	0.05	0.50 a	0.00	15 a	0.84	3.80
F	0.95 a	0.30 b	0.20	26 a	0.66	5.42
Minimum	0.00	0.00	0.00	0	0.10	0.83
Maximum	1.00	0.80	1.05	38	0.98	7.47
Mean	0.36	0.26	0.34	17	0.71	2.86

N/D: Not Defined

Significant at a: 0.005; b: 0.01; c: 0.025 level (1-tailed t-test [since parameters must be ≥ 0]).

feedbacks or dynamics. The improvement is consistent as well with the results of Bowman's (1963) application of similar rules to inventory management data for actual firms.

6. Misperceptions of Feedback

The results strongly support the hypothesis that subjects use the proposed heuristic to manage their inventories. Several issues may now be addressed. What do the estimated parameters reveal about the causes of the severely dysfunctional performance of the subjects? To what causes do subjects attribute the dynamics they experience, and how do these attributions affect the potential for learning? Finally, why do subjects use a rule that produces such poor results? The results reveal several distinct misperceptions of the feedback structure of the simulated environment. These misperceptions are responsible for the poor performance of the subjects.

Anchoring in the Choice of the Desired Stock

How do subjects select the desired stock? Because the complexity of the system and limited time available make calculation of optimal inventory levels infeasible, it is hypothesized that the subjects' choice of S^* is anchored to the initial level of 12 units. Since $S' = S^* + \beta SL^*$, S^* and SL^* may be estimated by regression of the estimated values of β on S':

$$S' = \underset{(6.9)}{13.9} + \beta * \underset{(2.8)}{8.4}, \qquad N = 40, \qquad R^2 = 0.09. \tag{14}$$

The low R^2 indicates, as one might expect, that individual differences in S^* and SL^* account for most of the variance in S'. The estimated value of SL^*, significant at the 10% level, is considered below. The estimated value of the desired stock S^*, that is, the value of S' when $\beta = 0$, is not significantly different from the initial level of 12 units. It appears that in the absence of a calculus to determine optimal inventories, subjects strongly anchor desired stocks on their initial level.

Misperception of Time Lags

To understand the source of the oscillations it is necessary to consider how the subjects dealt with the long time lags between placing and receiving orders—the supply line. The results show that most subjects failed to account adequately for the supply line. The evidence takes two forms. First, the small estimate of SL^* found in equation (14) indicates that the subjects underestimated the lag between placing and receiving orders. To ensure an appropriate acquisition rate the supply line must be proportional to the lag in acquiring beer (equation (7)). The acquisition lag is never less than 4 weeks (3 for the factory). Even if subjects' expectations of demand (and thus desired throughput) remained at the initial level of 4, the required supply line would be 16 cases, far greater than the estimated value of 8.4 cases. Thus it appears that subjects failed to allow for sufficient beer in the pipeline to achieve their desired inventory level.

More significant is the extent to which subjects responded to the supply line itself, as indicated by the estimated values of β. The optimal value of β is unity: subjects should fully account for the goods in the supply line to prevent overordering. But the mean value of β is just 0.34; only five subjects (11%) accounted for more than two-thirds of the supply line. The result is overordering and instability. For example, consider the Grizzly factory (Figure 4; $R^2 = 0.75$). As in most trials, the distributor begins to place substantially higher orders around week 15. These orders deplete the factory's inventory and build up a backlog of unfilled orders, encouraging the factory to boost orders. However, α_S for the Grizzly factory is 0.65 while $\beta = 0$, meaning the subject ordered two-thirds of the discrepancy between S' and S each period, and completely ignored the supply line.

Since the factory's supply line is three weeks long, the subject orders two-thirds of the stock shortfall for three successive weeks before receiving any of these new orders, overordering by a factor of two. Thus factory orders reach a peak of 50 units in weeks 18 and 19, coincident with the largest backlog. Inventory then rises toward the desired level and the subject cuts orders back. But the orders already in the pipeline continue to arrive, ultimately swelling inventory to a peak of 69 units. Because the distributor also acquired excess inventory (the distributor's $\beta = 0.25$), the factory finds incoming orders plummet to an average of just 5 cases per week after week 25, and ends the trial with high inventory, no way to unload it, and considerable frustration. The factory's ordering policy significantly amplifies the distributor's orders: incoming orders rise from 4 to 20 units; the factory responds by raising orders from 4 to 50 units, an amplification factor of 290%. By ignoring the supply line the factory's ordering policy is highly destabilizing.

In contrast, consider the Suds factory (Figure 4, $R^2 = 0.95$). Here $\beta \approx 1$ while $\alpha_S = 0.35$, indicating the subject fully accounts for the supply line and seeks to correct 35% of any inventory discrepancy each period. Because the Suds factory accounted for the supply line, orders peak and fall *before* the backlog reaches its maximum since the subject realized that sufficient orders to correct the problem were already in the pipeline. The Suds factory actually stabilizes the system: the amplification factor is 85%, meaning the factory's ordering policy attenuates demand shocks rather than exacerbating them.

"Open-Loop" Explanations of Dynamics

At the end of the game subjects are debriefed. Emotions run high. The majority express frustration at their inability to control the system. Many report feelings of helplessness—they feel themselves to be at the mercy of forces outside their control. Subjects are then asked to sketch their best estimate of the pattern of customer demand, that is, the contents of the customer order deck. Only the retailers have direct knowledge of that demand. Figure 5 shows a typical set of responses. Invariably the majority of subjects judge that customer demand was oscillatory, first rising from the initial level of 4 cases per week to a peak anywhere from 12 to 40 cases, and then dropping to the neighborhood of 0 to 12 cases per week. Factories and distributors tend to draw the largest excursion; wholesalers tend to draw smaller fluctuations. Only a small fraction suggest that customer demand was essentially constant. It may seem obvious that subjects' judgments of customer demand reflect their experiences during the game: after all, customer demand in reality does fluctuate. Yet these beliefs are revealing. Most subjects attribute the cause of the dynamics they experienced to external events. Most blame their own poor performance on what they see as a perverse pattern of customer demand: the customers increased their demand, encouraging them to order additional beer, but suddenly stopped ordering

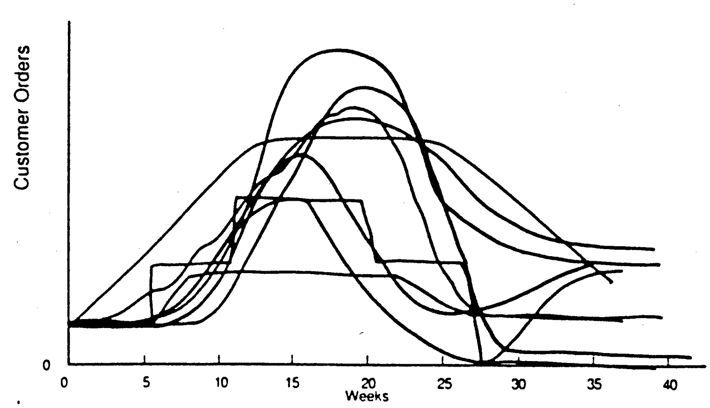

FIGURE 5. Typical Sample of Subjects' Post-Play Judgments of Customer Orders. Compare against actual customer orders (Figure 3).

just when the tap began to flow. Many participants are quite shocked when the actual pattern of customer orders is revealed; some voice strong disbelief. Few ever suggest that their own decisions were the cause of the behavior they experienced. Fewer still explain the pattern of oscillation in terms of the feedback structure, time delays, or stock and flow structure of the game.

Most subjects attribute the dynamics to external variables which they believe to be closely correlated in time and space with the phenomenon to be explained. These explanations reflect an 'open-loop' conception of the origin of dynamics, as opposed to a mode of explanation in which change is seen as arising from the endogenous interactions of decision makers with their environment. Learning from experience may be hindered by such misperceptions of the origins of dynamic behavior. When asked how they could improve their performance, many call for better forecasts of customer demand. The erroneous open-loop attribution of dynamics to exogenous events thus draws subjects' efforts to learn away from the high leverage point in the system (the stock management policy) and towards efforts to anticipate and react to external shocks. While better forecasts are likely to help, the key to improved performance lies within the policy individuals use to manage the system and not in the external environment. Even a perfect forecast will not prevent a manager who ignores the supply line from overordering.

7. Discussion and Conclusions

The experiment, despite its rich feedback structure, is vastly simplified compared to the real world. To what extent do the experimental conditions and results apply? First, would subjects' behavior differ if customer demand followed a more realistic pattern, e.g. noise or seasonality? The order decisions of many subjects were in fact noisy and cyclic (Figure 4). Therefore subjects upstream of these noisy individuals did in fact experience realistic demands. The behavior of these subjects is not statistically different from that of the retailers, indicating that the use of a step input does not reduce the generality of the results.

More fundamentally, are the main features of the experimental behavior observed in real production-distribution systems? It has long been recognized that production-distribution networks in the real economy exhibit the three aggregate behaviors generated in the experiment, i.e. oscillation, amplification from retail sales to primary production, and phase lag (T. Mitchell 1923, Hansen 1951, W. Mitchell 1971, Zarnowitz 1973). Is it plausible that managers in the real economy fall victim to the same misperceptions of feedback which plague subjects of the experiment? After all, in reality managers have access to more information than is available in the experiment. More time is available to gather intelligence and deliberate. Decision aids may be used. On the other hand information in the real world is often out of date, noisy, contradictory and ambiguous. Managers struggle to balance competing demands on their time and must make many additional decisions besides the quantity of goods to order. Consultants and models are subject to many of the same cognitive, informational, and temporal limitations, and there is no accepted calculus for integrating numerous and possibly conflicting positions and information sources.

The hypothesis that managers in real stock management contexts use a rule like the proposed anchoring and adjustment heuristic does not require equivalence of the decision-making tasks but only the weaker condition that in both cases the determination of optimal quantities exceeds the abilities of the decision makers. The virtue of the rule is its simplicity. It requires no knowledge of the dynamics or general equilibrium of the system. It is self-correcting—the feedback structure of the rule ensures that forecast errors, changes in the structure of the environment, and even self-generated overreactions can eventually be corrected. The benchmark costs (Table 2) show the rule can, with reasonable parameters, produce excellent results. As argued in Sterman (1987a), the decision rule characterizes actual decisions well because it captures the essential attributes of any minimally sensible stock management procedure. These are replacement of expected losses, correction of discrepancies between the desired and actual stock, and an accounting for the supply line of unfilled orders.

Of course individual managers do not ignore the goods they have on order. The problem in the real economy is one of aggregation. There are many examples of stock management situations in which the aggregate supply line is distributed among individual competitors and largely unknown to each. It is interesting to note that many of the markets most prone to instability such as agricultural commodities, commercial construction, machine tools, electronic components, and other durable goods are characterized by both significant delays in bringing investments to fruition and imperfect knowledge of the plans, commitments, and pending investments of the participants (Meadows 1970, Hoyt 1933, Commodity Research Bureau, various years). Verification of the supply line hypothesis requires further empirical work focussed not only on the decision processes of individual firms but also on the availability, timeliness, salience, and perceived accuracy of supply line information.

The robustness of the stock management heuristic is illuminating here. An earlier experiment tested the heuristic in a macroeconomic context (Sterman 1987a, 1989). Subjects were responsible for capital investment decisions in a simulated multiplier-accelerator economy. In contrast to the beer game, with its complex structure, multiple players, and time pressure, the macroeconomic system was rather simple. Perfect information was available to the subjects. There were no other participants to consider. The cost function was symmetric. There was no time limit. Yet as in the beer game, the results strongly supported the proposed rule. The rule explained an average of 85% of the variance of the subjects' decisions, and the estimated parameters were generally highly significant. As in the beer game, performance was decidedly suboptimal. Subjects produced large amplitude cycles in response to nonoscillatory inputs. The same misperceptions of feedback were apparent. In particular, subjects were insensitive to the presence of feedback from their decisions to the environment, underestimated the time lag between action and response, and failed to account for the supply line.

Though the stock-management task investigated here has wide applicability, there are many dynamic decision-making tasks which cannot be described by that framework (e.g. price-setting behavior). However, the results suggest the method used here may be helpful in explaining how unintended and dysfunctional results may be produced by apparently reasonable decision processes in diverse systems (e.g. Hall's account (1976, 1984) of the *Saturday Evening Post* and other organizations). Morecroft (1985) suggests the use of simulation to test the intended rationality of the decision rules in simulation models. The experimental approach used here allows direct investigation of the decision processes of real managers, and provides a technique to relate these decision rules to performance. Normative use of the techniques appears also holds some promise.[7]

Future work should apply the experimental method used here to other dynamic decision tasks and should consider the processes by which the parameters of the heuristics are modified or the heuristics themselves revised or replaced by learning and the selective pressures of the market. Tversky and Kahneman (1987) and Hogarth (1981) have stressed ways in which inadequate outcome feedback may hinder learning and efficiency. The

[7] In a study in progress, a similar game has been developed for an insurance company. Like the beer game, it appears that similar underperformance and misperceptions arise. After estimating the parameters of the managers' decision rules, the sources of poor performance will be discussed in training sessions. It is hoped that such training will help managers develop more appropriate heuristics by improving their mental models of the feedback environment.

results here suggest that outcome feedback alone is not sufficient: by attributing the source of change to external factors, people's mental models lead them away from the true source of difficulty. Efforts to improve performance may therefore have little leverage and experience may not lead rapidly to improved mental models, allowing dysfunctional performance to persist.

These results reinforce and extend prior work in dynamic decision-making (Brehmer 1987, Hogarth 1981, Kleinmuntz 1985, MacKinnon and Wearing 1985, Remus 1978). The efficacy and robustness of decision strategies lies not only in the availability of *outcome* feedback, but depends crucially on the nature of the *action* feedback between decisions and changes in the environment which condition future decisions. A heuristic may produce stable behavior in one setting and oscillation in another solely as a function of the feedback structure in which it is embedded. That structure consists of the stock and flow structure, information networks, time delays, and nonlinearities which characterize the organization. The magnitude of the oscillations despite a virtually constant external environment suggests the powerful role of action feedback in the genesis of dynamics. Further, the qualitative behavior of the different teams is strikingly similar despite wide variation in individual responses (as represented by the diverse parameters which characterize different subjects). As a result, the aggregate dynamics of an organization may be relatively insensitive to the decision processes of the individual agents, suggesting the importance in both descriptive and normative work of research methods which integrate individual decision-making with theories of feedback structure and dynamics. In that spirit the results show how experimental methods may be coupled with simulation to form a useful part of the "apparatus for moving from the level of the individual actor to the behavior of the system," ultimately yielding testable theories to explain the endogenous generation of macrobehavior from the microstructure of human systems.[8]

[8] The comments of John Carroll, Richard Day, James Hines, Robin Hogarth, Don Kleinmuntz, Robert Winkler, and anonymous referees are gratefully acknowledged. Daniel Ryu provided invaluable assistance.

References

BOWMAN, E. H., "Consistency and Optimality in Managerial Decision Making," *Management Sci*., 9 (1963), 310–321.

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

CAMERER, C., "General Conditions for the Success of Bootstrapping Models," *Organizational Behavior and Human Performance*, 27 (1981), 411–422.

COLEMAN, J. S., "Psychological Structure and Social Structure in Economic Models," in R. Hogarth and M. Reder (Eds.), *Rational Choice: The Contrast Between Economics and Psychology*, University of Chicago Press, Chicago, 1987.

COMMODITY RESEARCH BUREAU, *Commodity Year Book*, Commodity Research Bureau, Inc., New York, various years.

CYERT, R. AND J. MARCH, *A Behavioral Theory of the Firm*, Prentice Hall, Englewood Cliffs, NJ, 1963.

DAVIS, H. L., S. J. HOCH AND E. K. EASTON-RAGSDALE, "An Anchoring and Adjustment Model of Spousal Predictions," *J. Consumer Res*., 13 (1986), 25–37.

EINHORN, H. J. AND R. M. HOGARTH, "Ambiguity and Uncertainty in Probabilistic Inference," *Psychological Rev*., 92 (1985), 433–461.

———, D. KLEINMUNTZ AND B. KLEINMUNTZ, "Linear Regression and Process-Tracing Models of Judgment," *Psychological Rev*., 86 (1979), 465–485.

FRANKEL, J. A. AND K. A. FROOT, "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *Amer. Economic Rev*., 77 (1987), 133–153.

GRETHER, D. AND C. PLOTT, "Economic Theory of Choice and the Preference Reversal Phenomenon," *Amer. Economic Rev*., 69 (1979) 623–638.

HALL, R. I., "The Natural Logic of Management Policy Making: Its Implications for the Survival of an Organization," *Management Sci*., 30 (1984), 905–927.

———, "A System Pathology of an Organization: The Rise and Fall of the Old Saturday Evening Post," *Admin. Sci. Quart*., 21 (1976), 185–211.

HANSEN, A. H., *Business Cycles and National Income*, W. W. Norton, New York, 1951.
HINES, J. H., "A Behavioral Theory of Interest Rate Mechanics," *Essays in Behavioral Economic Modeling*, PhD dissertation, Sloan School of Management, MIT, 1987.
HOGARTH, R. M., "Beyond Discrete Biases: Functional and Dysfunctional Aspects of Judgmental Heuristics," *Psychological Bulletin*, 90 (1981), 197–217.
——— AND M. W. REDER (Eds.), *Rational Choice: The Contrast Between Economics and Psychology*, University of Chicago Press, Chicago, 1987.
HOYT, H. *One Hundred Years of Land Values in Chicago*, University of Chicago Press, Chicago, 1933.
JOHNSON, E. J. AND D. A. SCHKADE, "Heuristics and Bias in Utility Assessment," unpublished manuscript, Wharton School, University of Pennsylvania, Philadelphia, 1987.
JUDGE ET AL., "*The Theory and Practice of Econometrics*, Wiley, New York, 1980.
KAHNEMAN, D. AND A. TVERSKY, "The Simulation Heuristic," in Kahneman, D. et al., *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, 1982.
———, P. SLOVIC AND A. TVERSKY, *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, 1982.
KLEINMUNTZ, D. N., "Cognitive Heuristics and Feedback in a Dynamic Decision Environment," *Management Sci.*, 31, 6 (1985), 680–702.
MACKINNON, A. AND A. WEARING, "Systems Analysis and Dynamic Decision Making," *Acta Psychologica*, 58 (1985), 159–172.
MEADOWS, D. L., *Dynamics of Commodity Production Cycles*, MIT Press, Cambridge, MA, 1970.
MITCHELL, T. W., "Competitive Illusion as a Cause of Business Cycles," *Quart. J. Economics*, 38 (1923), 631–652.
MITCHELL, W. C., *Business Cycles and their Causes*, Univ. of California Press, Berkeley, 1971.
MORECROFT, J., "System Dynamics: Portraying Bounded Rationality," *Omega*, 11 (1983), 131–142.
———, "Rationality in the Analysis of Behavioral Simulation Models," *Management Sci.*, 31 (1985), 900–916.
PLOTT, C. R., "Laboratory Experiments in Economics: The Implications of Posted Price Institutions," *Science*, 232 (9 May 1986), 732–738.
REMUS, W. E., "Testing Bowman's Managerial Coefficient Theory Using a Competitive Gaming Environment," *Management Sci.*, 24 (1978), 827–835.
SIMON, H. A., "Rational Decisionmaking in Business Organizations," *Amer. Economic Rev.*, 69 (1979), 493–513.
SIMON, H. A., *Models of Bounded Rationality*. The MIT Press, Cambridge, MA, 1982.
———, "The Behavioral and Rational Foundations of Economic Dynamics," *J. Economic Behavior and Organization*, 5 (1984), 35–55.
SLOVIC, P. AND S. LICHTENSTEIN, "Preference Reversals: A Broader Perspective," *Amer. Economic Rev.*, 73 (1983), 596–605.
SMITH, V. L., "Experimental Methods in the Political Economy of Exchange," *Science*, 234 (10 October 1986), 167–173.
———, "Microeconomic Systems as an Experimental Science," *Amer. Economic Rev.*, 72 (1982), 923–955.
STERMAN, J. D., "A Behavioral Model of the Economic Long Wave," *J. Economic Behavior and Organization*, 6 (1985), 17–53.
———, "Testing Behavioral Simulation Models by Direct Experiment," *Management Sci.*, 33 (1987a), 1572–1592.
———, "Expectation Formation in Behavioral Simulation Models," *Behavioral Sci.*, 32 (1987b), 190–211.
———, "Misperceptions of Feedback in Dynamic Decisionmaking," *Organizational Behavior and Human Decision Processes*, forthcoming 1989.
TVERSKY, A. AND D. KAHNEMAN, "Judgment Under Uncertainty: Heuristics and Biases," *Science*, 185 (27 September 1974), 1124–1131.
——— AND ———, "The Framing of Decisions and the Psychology of Choice," *Science*, 211 (30 January 1981), 453–458.
——— AND ———, "Rational Choice and the Framing of Decisions," in R. Hogarth and M. Reder (Eds.), *Rational Choice: The Contrast Between Economics and Psychology*, University of Chicago Press, Chicago, 1987.
ZARNOWITZ, V., *Orders, Production, and Investment—A Cyclical and Structural Analysis*, National Bureau of Economic Research, New York, 1973.