Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decisionmaking Experiment

John D. Sterman
Associate Professor
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02139

ABSTRACT

Studies in the psychology of individual choice have identified numerous cognitive, informational, temporal, and other limitations which bound human rationality, often producing systematic errors and biases in judgment and choice. Yet for the most part models of aggregate phenomena in management science and economics have not adopted postulates of human behavior consistent with such micro-empirical knowledge of individual decisionmaking. One reason 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 macro-dynamics from microstructure in a common and important managerial context. Subjects play the role of managers in a simulated inventory management system, the "Beer Distribution Game". The simulated environment 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. Subjects generate large amplitude oscillations with stable phase and gain relationships among the variables. An anchoring and adjustment heuristic for stock management is proposed as a model of the subject's decision process. The parameters of the rule are estimated and the rule is shown to explain the subjects' behavior well. Analysis shows the subjects fall victim to several 'misperceptions of feedback' identified in prior experimental studies of dynamic decisionmaking. Specifically, they fail to account for control actions which have been initiated but not yet had their effect. More subtle, subjects are insensitive to the presence of feedback from their decisions to the environment and attribute the dynamics to exogenous variables, leading their normative efforts away from the source of difficulty. The experimental results are related to prior tests of the proposed heuristic and the generality of the results is considered. Finally implications for behavioral theories of aggregate social and economic dynamics are explored.

INTRODUCTION

Economics and psychology, despite their common focus on human behavior, have been locked in battle for much of the past century. The battle centers on the assumptions about decisionmaking behavior upon which theories of choice are to be based. At the risk of oversimplification, the two positions can be characterized as follows.[1] Economists favor theories based on axioms of rational choice. Decisionmaking behavior is assumed to be rational and consistent. Agents maximize utility or profits, and the information required to do so is either freely available or optimally purchased. In the most extreme form, exemplified today by rational expectations models, agents have perfect models of the economy and never systematically err. In contrast, psychologists have arrayed on the battlefield a formidable host of experimental results documenting departures from optimal behavior in a wide variety of decisionmaking tasks. Rationality is bounded by limitations of information, time, and cognitive capability (Simon 1979). Preferences are frequently intransitive and are shaped by alternate modes of elicitation and framing (Slovic and Lichtenstein 1983). Individuals are inconsistent and can often be outperformed by simple models (Kleinmuntz 1985, Hogarth and Makridakis 1981a, Goldberg 1976, Dawes 1971, Bowman 1963). Experiments have identified numerous heuristics commonly used in prediction and decisionmaking and a wide array of systematic errors to which these heuristics are prone (Tversky and Kahneman 1974, Kahneman, Slovic, & Tversky 1982, Hogarth and Makridakis 1981b).

The broad gulf between the perspectives elicits diverse reactions. Zeckhauser (1986) argues that the debate has many of the properties of a Kuhnian paradigm conflict: each side can score points on their own court at will, but few on either side are convinced to change. Others, notably Leontief (1971), Phelps-Brown (1972) and Simon (1984) 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, 40).

Much of the empirical work in experimental economics and the psychology of choice has generated just such microlevel data (Einhorn and Hogarth 1981, Plott 1986, Smith 1986). But the focus of much research in behavioral decision theory on individual behavior in static and discrete tasks has limited the penetration of psychological perspectives in theories of aggregate dynamics such as the behavior of markets, firms, and other economic systems. 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 a subject'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. As a result it has been difficult for behavioral decision theory to make much headway in analyzing the dynamics of aggregate organizations such as a firm or industry. Coleman (1986) argues that the greatest progress in bridging the gulf 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 and important managerial context. In the experiment subjects play the role of managers in a simulated industrial production and distribution system, the "Beer Distribution Game". The decisionmaking 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. Subjects generate large amplitude oscillations with stable phase and gain relationships among the variables. An anchoring and adjustment heuristic for stock management is proposed as a model of the subject's decision process. The parameters of the rule are estimated and the rule is shown to explain the subjects' behavior well. Analysis of the results shows the subjects of the experiment 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. Most important, subjects were insensitive to the presence of feedback from their decisions to the environment. The majority attribute the dynamics to external events, when in fact the dynamics they experience are internally generated by their own actions. These misperceptions are shown to be responsible for the poor performance of the subjects. Further, the subjects' open-loop mental model, in which dynamics arise from exogenous forces, is hypothesized to hinder learning and retard evolution towards greater efficiency.

The experimental results are related to prior tests of the proposed heuristic (Sterman 1987a, 1987b) and the generality of the results is considered. It is shown that the same regularities evident in the subjects' behavior appear in real-world production-distribution systems. Finally implications for behavioral theories of aggregate social and economic dynamics are explored.

THE STOCK MANAGEMENT PROBLEM

One of the most common dynamic decisionmaking tasks is the regulation of a stock or system state. In such a problem, the manager seeks to maintain a quantity at a particular target level, or at least within an acceptable range. Typically the stock cannot be controlled directly but rather is influenced by altering the rates of flow which accumulate into and out of the stock. The manager must set the inflow rate so as to compensate for losses from the stock and to counteract disturbances which push the stock away from its desired value. Frequently there are lags between the initiation of a control action and its effect on the stock, and/or lags between a change in the stock and the perception of that change by the decisionmaker. 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 from the micro to the macro. 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 piling up. 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 so as to provide sufficient credit for economic growth while avoiding inflation, 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 of interest S is the accumulation of the acquisition rate A less the loss rate L:

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

Losses from the stock must depend on the stock itself, and may also depend on other endogenous variables **X** and exogenous variables **U**:[2]

$$L_t = f_L(S_t, X_t, U_t). \tag{2}$$

The acquisition rate will depend 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 be a function of the supply line itself and the other endogenous and exogenous variables:

$$A_t = f_A(SL_t, \lambda_t). \tag{3}$$

The supply line is simply the accumulation of the orders which have been placed O less those which have been delivered:

$$SL_t = \int_{t_0}^{t} (O_\tau - A_\tau)d\tau + SL_{t_0} \tag{4}$$

The structure represented by figure 1 and eq. (1)-(4) is quite general. There is no presumption that the functions governing losses and the acquisition lag are linear. 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's task is to choose the order rate over time so as to keep the stock close to a target.[3] It is interesting to note that the characteristic behavior modes of many of these systems include oscillation and instability.

There are two extreme approaches to modeling the decision process used to determine orders. At one extreme, one may assume that the manager chooses the path of orders optimally with respect to some objective function. At the other extreme, one may assume the decisionmaker is random, i.e. that there is no control at all. The model proposed here is an intermediate one. It assumes that managers are unable to optimize and instead utilize a heuristic which is locally rational. The proposed heuristic 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. Local rationality in the context of simulation models is discussed by Morecroft 1983, 1985 and Sterman 1985, 1987a.

The proposed decision rule thus utilizes information locally available to the decisionmaker, and does not presume the manager has global knowledge of the structure of the system. The generic decision rule recognizes three motives for ordering:

> Order enough 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.

1. *Replacement of losses.* The replacement motive is straightforward. In equilibrium, when the desired and actual stock are equal, the manager must continue to order enough to replace ongoing losses. Losses may arise from usage (as in a raw material inventory) or decay (as in the depreciation of plant and equipment). Failure to replace losses would cause the stock to fall below the desired level, creating steady-state error. Accurate forecasts of losses allow replacement alone to maintain the stock close to its desired value.

2. *Stock adjustment.* The possibility of forecasting errors or changes in the desired stock demands a mechanism to adjust orders above or below replacement. Orders to reduce the discrepancy between the desired and actual stock form a negative feedback loop which regulates the stock (shown in the bottom part of figure 1). Any rule which fails to compensate for discrepancies between the desired and actual stock fails to control the stock at all. Such a rule could not respond to a change in the desired stock, nor restore the stock to the desired value if displaced. The stock would follow a random walk as the system is bombarded by shocks.

3. *Supply line adjustment.* Delays between the initiation and impact of control actions give stock management systems significant inertia and should be accounted for by managers to ensure a stable response to shocks. The importance of the supply line adjustment can be illustrated with two simple examples. Consider first hammering a nail into a board, the classic example used by Miller, Galanter, and Pribram (1960) to illustrate the use of feedback in their concept of the TOTE unit (Test-Operate-Test-Exit) as a structure for managing systems. "In a TOTE unit, one Tests to see if a goal is met, Operates to approach the goal, Tests again, and Exits from the loop when the goal is reached" (Richardson 1984, 292). The nail-hammer system is a simple stock-management situation in which the system state to be managed is the distance between the nail head and the surface of the board. The desired state is to have the nail head flush with the board. The following decision rule for hammering will work well:

```
Routine "Hammer"
  TEST: is nail_height>0?
  if yes, then OPERATE: hammer!
    goto TEST
  if no, then EXIT.
```

The decision rule implements a simple negative feedback loop whose goal is to reduce the distance between nail head and board to zero. Note the simplicity of the feedback system: there are no losses or external disturbances which influence the state of affairs (the nail does not pull itself out of the board), and there are no significant time lags between striking a blow and the distance remaining, or between a change in the distance remaining and the perception of that change.

Now apply the same logic to the problem of ordering dinner in a restaurant. The desired stock is a full stomach, the actual stock is initially an empty one. Dinner is ordered in response to the discrepancy between the desired and actual stock:

```
Routine "Dine Out"
  TEST: is hunger>desired hunger?
  if yes, then OPERATE: order dinner
    goto TEST
  if no, then EXIT.
```

Using this rule one would order another meal each time the waiter passed by (since the discrepancy between desired and actual stock would still exist). You would only stop ordering when the first dinner was served. Soon your table would be piled high with redundant dinners. A stock management heuristic which fails to measure and respond to the supply line of unfilled orders is predisposed to instability.[4]

Despite its importance it is not obvious a priori that people actually do attend to the supply line. In many stock management situations the lag between action and response so short that the supply line can be effectively ignored, as in the case of the nail. In others information about the supply line is not available or salient, as in a decentralized market where each participant is unaware of the plans of the others. We might question the rationality of the person who inadvertently orders several dinners in a restaurant, but consider cooking dinner at home on an electric range. Who among us has never overcooked a meal by failing to account for the supply line of heat in the coils of the range which, even after the burner has been turned off, continues to heat the pot? Whether managers account for the supply line is an empirical question in any particular situation.

Formalizing the Heuristic

The following equations formalize the ordering heuristic proposed above. First, orders in most real life situations must be nonnegative,

$$O_t = MAX(0, IO_t) \tag{5}$$

where IO is the indicated order rate, the rate indicated by other pressures.[5]

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 their effects by what Kahneman and Tversky (1982) call 'mental simulation.' Anchoring and adjustment has been shown to apply to a wide variety of decisionmaking tasks (Einhorn and Hogarth 1985, Davis, Hoch, and Ragsdale 1986, Johnson and Schkade 1987, Hines 1987, Lopes 1981). Here the anchor is the expected loss rate L^e. Adjustments are then made to correct discrepancies between the desired and actual stock AS, and between the desired and actual supply line ASL:

$$IO_t = L^e{}_t + AS_t + ASL_t. \tag{6}$$

The expected loss rate may be formed in various ways. Common assumptions in economics and management science include static expectations $L^e{}_t = L^*$ (a constant or equilibrium value), regressive expectations $L^e{}_t = \gamma L_{t-1} + (1-\gamma)L^*$, $0 \le \gamma \le 1$, adaptive expectations $L^e{}_t = \theta L_{t-1} + (1-\theta)L^e{}_{t-1}$, $0 \le \theta \le 1$, and extrapolative expectations, $\Delta L^e{}_t = \Sigma \omega_i * \Delta L_{t-i}$, where Δ is the first difference operator and $\omega_i > 0$.

The adjustment for the stock AS creates the chief negative feedback loop which regulates the stock. The proposed heuristic assumes for simplicity that the adjustment is linear in the discrepancy between the desired stock S* and the actual stock:

$$AS_t = \alpha_S(S^*{}_t - S_t), \tag{7}$$

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

$$ASL_t = \alpha_{SL}(SL^*{}_t - SL_t), \tag{8}$$

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:

$$SL^*{}_t = \lambda^e{}_t * \Phi^*{}_t. \tag{9}$$

The longer the expected delay in acquiring goods or the larger the throughput desired, the larger the quantity on order 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 lag in acquiring orders. The supply line adjustment thus avoids overordering (as in the restaurant example) and also compensates for changes in the acquisition lag. For example if the acquisition lag doubled 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 λ^e and Φ^*, ranging from constants through sophisticated forecasts.[6]

In terms of the anchoring and adjustment heuristic, the expected loss forms 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 decisionmaker. 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 as given in equations (5)-(9). Rather, pressures arising from the discrepancies between desired and actual stock and desired and actual supply line cause managers to adjust the order rate above or below the level which would maintain the status quo. The adjustment parameters α_S and α_{SL} reflect the manager's response to disequilibrium: large values indicate aggressive efforts to bring the stock and supply line in line with their desired levels, respectively; small values indicate a cautious approach, or less sensitivity to discrepancies between desired and actual stocks. The negative feedback loop structure of the rule reduces the sensitivity of the results to the adjustment parameters: if the initial response to disequilibrium is insufficient, additional adjustments will be made until balance is restored; overcorrection will likewise ultimately itself be corrected. These self-correcting feedbacks allow the heuristic to be used in any stock management situation without detailed knowledge of its dynamics.

Prior Tests of the Proposed Heuristic

The proposed heuristic has a long history in economics and management science. Variants of the rule have been used in models of aggregate capital investment (e.g. Samuelson 1939, Hall and Jorgenson 1967) and production planning at the level of the firm (Holt, Modigliani, Muth, and Simon 1960, Forrester 1961), among others. However, these rules were not tested experimentally but were postulated ad hoc as 'reasonable' or justified as optimal under certain restricted conditions (e.g. quadratic costs). A recent experiment (Sterman 1987a, Sterman 1987b) tested the proposed rule in a macroeconomic context. Subjects were responsible for capital investment decisions in a simulated multiplier-accelerator economy. The results strongly supported the proposed rule. The rule explained an average of 85% of the variance of the subject's decisions, and the estimated parameters were generally highly significant. The performance of the subjects was decidedly suboptimal. Subjects produced large amplitude cycles in response to nonoscillatory inputs. The analysis revealed several misperceptions of feedback structure on the part of the subjects. 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 control actions which had been initiated but not yet had their effect.

While the macroeconomic experiment was suggestive several issues regarding the generality of the results remain. The experiment was a one-person game. Would subjects use a different heuristic in the presence of multiple players and the possibility for strategic behavior (gaming) thus created? The simulated economy was quite simple and the number of possible inputs to the decisions of the subjects was severely limited. How would a more complex feedback environment with many information sources influence stock management behavior? The "cover story" of the experiment was an aggregated macroeconomic setting. Would stock management behavior differ in a different task environment, specifically a task at the level of an individual firm? In short, does the stock management heuristic apply to other situations? Do the misperceptions of feedback structure identified in the earlier experiment arise in other stock management situations, or were they artifacts of the experiment?

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 in a simplified fashion the production and distribution of beer (figure 2). Orders for and cases of beer are represented by markers and pennies which are physically manipulated by the players as the game proceeds. Each brewery consists of four sectors: retailer, wholesaler, distributor, and factory (R, W, D, F). One person manages each sector. Customer demand is represented on a deck of cards. 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 over the length of the game. Costs are incurred at each link of the distribution chain as follows. Inventory holding costs are $.50 per case per week, and stockout costs (costs for having a backlog of unfilled orders) are $1.00 per case per week.

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 and satisfying customer demand. Inventory cannot be controlled directly but must be ordered. The lag in receiving beer is potentially variable: if the wholesaler has beer sufficient to cover the retailer's orders the retailer will receive the beer desired after three weeks. But if the wholesaler has run out, the retailer must wait until the wholesaler can receive additional beer from the distributor. Only the factory, the primary producer, faces a constant delay in acquiring inventory (there is no limit to the production capacity of the factory).

Experimental Protocol

The game is initialized in equilibrium. Each inventory contains 12 cases (pennies). Initial equilibrium throughput is four cases per week. Each shipping and production delay thus contains four cases, and each order slip reads four. Customer demand is initially four cases per week. To disturb the system customer demand increases to eight cases per week in week 5 and remains at that level thereafter (figure 3). The step input is used rather than, say, a more realistic pattern with seasonality, trends, or noise to simplify the analysis. The step creates a disequilibrium disturbance to which the subjects must react, while facilitating subsequent analysis.

A typical session would involve from three to eight teams of four players. Subjects are randomly assigned roles as retailer, wholesaler, etc. After description of the production and distribution system, each team briefly confers and selects a name for their brewery. The names are written on the blackboard. Each person is then asked to place $1 in a kitty to be wagered against the other teams.[7] The kitty goes to the team with the lowest total costs at the end of the game, winner take all. The cost function is explained and written on the blackboard, and the prohibition against communicating with teammates or other teams is announced. The game leader then explains the steps of the game (figure 2). 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, thus keeping the system in equilibrium. Beginning in week four the players are allowed to order any nonnegative quantity they wish. During the briefing as well as during play questions concerning rules, procedures, or interpretation are answered; questions concerning strategy or customer demand are not.

During play the game leader calls out the steps and writes the current week on the blackboard to keep each player and team in step. Occasionally players become confused and the facilitators will stop play until the problem is corrected. The subjects are told the game will run for fifty simulated weeks, but play is actually halted after about 36 weeks, thus avoiding 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 faces severe information limitations. Customer demand is not known to any of the subjects in advance. Each week, the retailer examines the top card on the customer order deck, fills those orders, and discards the card, face down. Thus retailers are the only subjects with direct knowledge of customer demand. Similarly, each person places their order slips face down in the 'orders placed' box. Thus each knows only the orders of their own customer, and these only after a delay of one week.

Subjects have good local 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 other players, either across or within a game. Even though the objective of each brewery is to minimize total costs, there is no process for the players to coordinate their decisions or jointly plan strategy. As in many real situations the problem of global optimization is factored into subgoals which are distributed throughout the organization. The players are, of course, sitting next to one another, so a certain amount of crosstalk and signalling is unavoidable. Each can readily look up and down the board and 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 players' outbursts may also convey information. Thus in contrast to the earlier experiment there are numerous sources of information which are potentially relevant and available to the subjects to assist them in making ordering decisions.

The sample

The game has been played hundreds of times with a wide range of people in many nations. The results reported here were drawn from four dozen games (192 subjects) collected over a period of four years. Since the records are kept manually by each player there are occasional accounting errors. A computer model of the game was used to test the records for consistency. Trials in which there were errors of more than a few cases per week for more than a few weeks in any of the four sectors were discarded from further analysis. Eleven games were retained, thus providing 44 subjects.[8] 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.

RESULTS

Comparison to optimal behavior

The complexity of the system (it is a 19th order nonlinear difference equation) renders calculation of the optimal behavior intractable. However, a benchmark for evaluating the performance of the subjects may be obtained through computer simulation. As implemented below, the proposed decision rule involves four parameters. The parameters which produce the minimum total costs were calculated by simulation of the game over the plausible parameter space.[9] The benchmark costs were computed subject to the same information limitations faced by the subjects. The minimum costs produced by the decision rule thus provide an upper bound for minimum costs. The benchmark costs are shown in table 2 compared to the 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. The subjects are clearly not producing behavior consistent with optimal management of the distribution system.

Behavioral Regularities

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 subjects' pattern of ordering reveals several regularities.

1. *Oscillation:* The trials are all characterized by instability and oscillation. The pattern of orders and of inventory is dominated by a large amplitude fluctuation with an average period of 21 weeks. Close examination of the behavior shows that 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 subjects tend to increase their orders. 'Effective inventory' is defined as inventory less any backlog of unfilled orders and generally become significantly negative, indicating the sectors have backlogs. The maximum backlog for the full sample averages 35 cases, and generally occurs between weeks 20 and 25. As additional product is brewed and shipped there is a surge in inventory levels, and inventory in many cases substantially overshoots its initial levels. The average peak inventory level is 40 cases, and occurs between weeks 25 and 30. Orders fall off rapidly as excess inventory builds up. Recalling that the cost function penalizes both backlogs and excess inventory it is clear that the large fluctuations of inventory over the cycle (the average

excursion of inventory is 75 cases) are responsible for the huge costs compared to the benchmark response.

2. *Amplification:* The amplitude of the excursion in orders increases steadily as one moves from customer to retailer to factory. The peak order rate at the factory level is on average more than double the peak order rate generated at the retail level. Likewise the variance of factory orders averages 5.5 times the variance of retail orders. 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%.[10] Amplification in inventory excursions is also apparent.[11]

3. *Phase lag:* The peak order rate tends to occur 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 decisionmaking and order delays from retailer to wholesaler and so on.[12]

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 decisionmaking behavior.

TESTING THE THEORY

To test the model the proposed decision rule must be adapted to the particular situation in the beer game and cast in a form suitable for estimation of the parameters. In the context of the beer game, the stock S corresponds to the 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 L^e, 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 corporate and economic systems, they are often a good model of the evolution of expectations in the aggregate (Sterman 1987c, Frankel and Froot 1987), and they are one of the simplest formulations for expectations flexible enough to adapt to a nonstationary process.

Each subject is free to determine the desired level of inventory S^* according to their own beliefs about how to minimize costs. Theory suggests the target inventory level should be chosen so as to minimize expected costs given the cost function and the expected variability of deliveries and incoming orders. However, the subjects do not have the time nor information to determine an optimal inventory level. The asymmetry of the cost function suggests desired inventory should be nonnegative. One might further hypothesize that in the absence of a procedure to calculate optimal inventory levels the subjects' choice of S^* would be strongly 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 have no direct way 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. It is therefore assumed that the desired supply line SL^* is constant, and thus SL^* becomes a parameter to be estimated.

The generic decision rule of eq. (5-9) then becomes:

$$O_t = MAX(0, IO_t), \tag{10}$$

$$IO_t = L^e_t + AS_t + ASL_t, \tag{11}$$

$$L^e_t = \theta L_{t-1} + (1-\theta)L^e_{t-1}, \quad 0 \le \theta \le 1, \tag{12}$$

$$AS_t = \alpha_S(S^* - S_t), \tag{13}$$

$$ASL_t = \alpha_{SL}(SL^* - SL_t), \tag{14}$$

where S^* and SL^* are constants. Defining $\beta = \alpha_{SL}/\alpha_S$ and $S' = S^* + \beta SL^*$ and collecting terms yields

$$IO_t = L^e_t + \alpha_S(S' - S_t - \beta SL_t). \tag{15}$$

Note that since S^*, SL^*, α_{SL} and α_S are all ≥ 0, $S' \ge 0$. Further, it is unlikely that subjects will place more emphasis on the supply line than on the inventory itself: the supply line does not directly enter the cost function nor is it as salient as the inventory. Therefore it is probable that $\alpha_{SL} \le \alpha_S$, meaning $0 \le \beta \le 1$. Thus β can be interpreted as the fraction of the supply line taken into account by the subjects. If $\beta = 1$, the subjects fully recognize the supply line and do not double order. If $\beta = 0$, orders placed are forgotten until they arrive, encouraging overordering and instability, as in the restaurant example.

The decision rule contains four parameters to be estimated (θ, α_S, S', and β) and is nonlinear. To estimate the parameters an additive disturbance term is assumed:

$$O_t = MAX(0, IO_t + \varepsilon_t), \quad \varepsilon_t \sim N(0, \sigma^2). \tag{16}$$

The disturbances ε are assumed to be independent, identical, and normally distributed. In this case, maximum likelihood estimates of the parameters may be found by minimizing the sum of the squared errors Σe_t^2. Estimates for each sector of each trial were found by grid search of the parameter space subject to the constraints $0 \le \theta \le 1$ and α_S, S', $\beta \ge 0$.[13] Independence and normality of the errors implies 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).[14]

Comparing simulated and experimental results

The estimated parameters are displayed in table 4 together with R^2 and root mean square errors between estimated and actual orders. The mean R^2 is 71% (median 76%); 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.[15]

As a further test of the proposed decision rule the game was simulated using the rule as specified in eq. (10-15) and the estimated parameters for each sector in each trial. 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 θ, α_S, S', and β 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):

$$\mathrm{Costs}_{i,j} = \underset{(16.7)}{1.11}*\mathrm{Simulated\ Costs}(\theta_j,\alpha_{Sj},S'_j,\beta_j)_i; \quad i=R,W,D,F; \quad j=1,\ldots,11 \qquad (17)$$

$N=44$, $R^2 = .40$.

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

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 introduced high-frequency noise by changing orders from week to week (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 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.

The results strongly support the hypothesis that subjects use the proposed anchoring and adjustment heuristic to manage their inventories. Approximately three quarters of the variance in actual orders is explained by the proposed rule, and the vast majority of the estimated parameters are highly significant. 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? And finally, if the rule produces such poor results, why is it used?

MISPERCEPTIONS OF FEEDBACK

The results reveal several distinct misperceptions of the feedback structure of the simulated environment. These misperceptions are directly responsible for the poor performance of the subjects.

Anchoring in the choice of the desired stock

The complexity of the system and limited time for decisions make calculation of optimal inventory levels infeasible. It was therefore hypothesized that the choice of the desired stock S* would be strongly anchored to the initial level of 12 units. Though S* is not estimated directly, the results do allow S* to be imputed. Recalling that $S'=S^*+\beta SL^*$ it is clear that S* and SL* can be estimated by regression of the estimated values of β on S'. The regression yields the expected relationship:

$$\underset{(6.9)}{S'= 13.9} + \underset{(2.8)}{\beta*8.4}, \qquad N=40, \quad R^2 = .09. \qquad (18)$$

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 inventory level of 12 units. As hypothesized, in the absence of a calculus to determine optimal inventory levels, subjects' choice of desired inventory levels appears to be strongly anchored to the initial inventory.

Misperception of time lags

To understand the source and magnitude of the oscillation it is necessary to consider the adjustment parameters α_s and β which govern the response to disequilibrium. The optimal adjustment parameters for the decision rule, as determined by simulation, are $\beta = 1$ and $\alpha_s = 1$: the supply line is fully accounted for and the discrepancy in the stock is corrected each period in full. Intuitively, a full accounting for the supply line prevents overordering, as in the restaurant example. And when the supply line is fully accounted for, the speed of adjustment can be increased without destabilizing the system.

Inspection of the results shows 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 (18) 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 (eq. (9)). The acquisition lag is never less than 4 weeks. 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 average value of β is just .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=.75). As in most of the 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 restore inventory by ordering additional units of beer. However, α_S for the Grizzly factory is .65 while $\beta = 0$, meaning the subject ordered two-thirds of any discrepancy between S′and S each period, and completely ignored the orders in the supply line. Since the factory's supply line is three weeks long, the subject orders two-thirds of the required amount for three successive weeks before receiving any of these new orders. After three weeks inventory rises toward the desired level and the subject cuts orders back. But the orders already in the pipeline continue to arrive, ultimately swelling inventory above desired levels by nearly a factor of three. Thus factory orders reach a peak of 50 units in weeks 18 and 19, coincident with the largest backlog. Factory inventory subsequently reaches a peak of 69 units, well in excess of reasonable coverage of either equilibrium or actual distributor demand. Because the Grizzly distributor also acquired excess inventory (the distributor's β=.25), distributor orders plummet to an average of just 5 cases per week after week 25, and the factory ends the trial with high inventory, no way to unload it, and considerable frustration. Note that the factory's ordering policy significantly amplifies the distributor's orders: distributor 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=.95). For this subject β=1.05 while α_S=.35, indicating the subject fully accounts for the supply line and seeks to correct 35% of any discrepancy between S′and S each period. Compared to Grizzly the Suds distributor is more extreme, increasing orders to a peak of 50 cases in week 20. Nevertheless the response of the Suds factory is more stable than that of Grizzly. Because the Suds factory accounted for the supply line orders peak and fall *before* the backlog reaches its maximum since the subject realizes that sufficient orders to correct the problem are already in the pipeline. The Suds factory actually stabilizes the system: the amplification factor is 85%, meaning the parameters which characterize the factory attenuate 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 considerable 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 forces. 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, then pulled the rug out 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.

The subjects exhibit a strong tendency to attribute behavior to external variables which they believe to be closely correlated in time and space with the phenomenon to be explained. Such 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 decisionmakers with their environment. Such misperception of the origins of dynamic behavior has implications for the possibilities of learning from experience. When asked how they could do better many argue that performance would be improved through better forecasting of customer demand. The erroneous open-loop attribution of dynamics to exogenous events thus draws normative efforts 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 results show clearly that the source of the dynamics and the ability to improve performance lie within the policy individual people 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.

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? This question has several components: are the main features of the experimental behavior (oscillation, amplification, phase lag) observed in real production-distribution systems? If so, to what extent are the proposed heuristic and specifically the misperceptions of feedback identified in the experiment responsible for that behavior? How robust is the proposed heuristic in the face of the differences in information availability, time, and other factors between the experiment and reality? These are empirical questions which can and should be investigated at the micro level of individual firms. Nevertheless, the experimental results are suggestive.

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 to primary production, and phase lag (T. Mitchell 1923, Hansen 1951, W. Mitchell 1971, Zarnowitz 1973). Figure 6 shows detrended data for production of consumer goods, intermediate goods, and primary materials in the U.S. from 1947 to 1987. Production at all three stages fluctuates significantly over the business cycle, cycle amplitude and coherency grow as one moves from retail sales to materials, and the expected phase lags are apparent as well (table 5).

How plausible is it that managers in the real economy use the proposed heuristic, and if they use it, 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 arrive at a decision. Decision aids may be used. On the other hand information in the real world is often out of date, noisy, contradictory and ambiguous. Managers have far more 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 the proposed anchoring and adjustment heuristic rather than optimizing does not require equivalence of the decisionmaking tasks but only the weaker condition that in both cases the determination of optimal quantities exceeds the abilities of the decisionmakers. The robustness of the proposed stock management heuristic is illuminating here. The decision rule has been shown to be an excellent model of behavior in two distinct experimental settings. In the macroeconomic experiment the dynamic structure of the system was rather simple. There were no other participants to consider and therefore no game-theoretic component to the decision task. Perfect information was available to the subject. The cost function was symmetric. There was no time limit for each decision. In contrast, the beer game is substantially more complex. The underlying dynamic system is high order, has multiple nonlinearities, and involves numerous time lags. There are multiple decisionmakers whose behavior should be taken into account. Local information is good and limited information about the other sectors is available. The cost function is asymmetric. Subjects must make their decisions under time pressure. Yet the same heuristic explains decisionmaking in both experiments with a high degree of accuracy. In both cases people appear to be insensitive to the feedback environment, differences in individual performance are closely related to differences in the parameters estimated for each subject, and the same misperceptions of feedback are documented in both.

If the rule is prone to such major misperceptions and produces such grossly dysfunctional performance, why is it used? The virtue of the rule is its simplicity. It requires no knowledge of the feedback structure 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 works 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. It does not follow from the generality of the rule, however, that it is so flexible that it can be made to work in any situation. The rule is clearly inconsistent with any decisionmaking strategy based on global optimization or rational expectations.

How plausible is it that firms in the real economy fail to adequately account for the supply line? It is not credible that individual managers forget that they have goods 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.

Though the stock-management task investigated here has wide applicability, there are many dynamic decisionmaking 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 to be a promising avenue for future work.[16]

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 (1986) and Hogarth (1981) have stressed ways in which inadequate outcome feedback may hinder learning and efficiency. The results here suggest that outcome feedback alone may not be sufficient: by attributing the source of change to external factors people's mental models lead them away from the true source of their poor performance. Efforts to improve performance may therefore have little leverage and additional experience may not lead rapidly to improved mental models, allowing dysfunctional performance to persist.

These results reinforce and extend prior work in dynamic decisionmaking (Hogarth 1981, Kleinmuntz 1985, Mackinnon and Wearing 1985, Remus 1978). Not only does the efficacy and robustness of particular decision strategies depend crucially on the availability and nature of outcome feedback, but on the nature of the *action* feedback between decisions and changes in the environment which condition future decisions. The same heuristic may produce stable behavior in one setting and oscillation in another solely as a function of the feedback structure in which that heuristic 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 produced by the actors 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 individuals' responses (as represented by the diverse parameters which characterize the subjects across positions and teams). 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 research of research methods which integrate individual decisionmaking 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.

ACKNOWLEDGMENTS

The comments of John Carroll, James Hines, Robin Hogarth, Charles Holt, Don Kleinmuntz, David Schkade and seminar participants at the University of Texas at Austin, the University of Chicago, and MIT are gratefully acknowledged. Dan Ryu provided invaluable research assistance.

NOTES

1. See Hogarth and Reder (1986) for a full exposition of positions on both sides.
2. For any real quantity the loss rate must approach zero as the stock is depleted. However, there is no presumption in eq. (2) that the loss process is linear, nor that losses are independent of the age distribution of individual units in the stock.
3. In particular situations the choice of the desired stock and the meaning of 'close' will be influenced by the loss function perceived by the manager and the manager's priors regarding the variability of the environment. These choices are not in general separable from the dynamics of the system.
4. The example is not intended to criticize the TOTE methodology but rather serves to illustrate the importance of accounting for the supply line. More general TOTEs could account for the delay between ordering and receiving dinner.
5. Order cancellations are sometimes possible and may exceed new orders in extreme conditions (e.g. the U.S. nuclear power industry in the 1970s). Since cancellations are likely to be subject to different costs and administrative procedures than new orders they should be represented separately as a distinct outflow from the supply line rather than as negative orders.
6. A common specification in dynamic models is $\lambda^e=\lambda$ and $\Phi^*= L^e$.
7. Protocols for experimental economics (e.g. Smith 1982) call for significant monetary rewards geared to performance in the task. However a number of experiments have shown performance is not significantly improved and may be worsened by increases in 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 (depending on the number of other teams). Though these are small rewards they serve to emphasize the goal of minimum *team* costs and appear to have a powerful motivating effect.
8. Analysis showed a slight tendency for the trials with the most extreme amplitude and highest costs to be most prone to accounting errors. Thus the final sample of eleven trials is biased slightly towards those who understood and performed best in the game. The effect is modest, however, and reinforces the conclusions drawn below regarding misperceptions of the feedback structure by the subjects.
9. To reduce the search space it was assumed that all four sectors were characterized by the same parameters. The optimal parameters are $\theta=0$, $\alpha_S=1$, $\beta=1$, and $S'=28$ (20 for the factory).
10. Amplification is a rough measure of the closed-loop gain of the system 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.
11. Note that the average period and excursion of factory inventory are somewhat less than that 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 behavior of the subjects illustrates the extent to which the feedback structure of the task shapes the behavior of the subjects.
12. There is no apparent lag between retailer and wholesaler or between distributor and factory, perhaps indicating subjects' use of information outside of their own sector: e.g. the factory may look at the distributor's inventory when ordering.
13. The parameters (θ, α_S, S', and β) were estimated to the nearest .1, .05, 1, and .05 units, respectively. The search was carried out over a sufficiently large range to ensure capturing the global minimum of Σe_t^2. The data and computer programs are available from

the author upon request.

14. Note, however, that the ordering function does not contain a regression constant. Therefore the residuals will not, in general, 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 goodness of fit. The alternative $R^2 = r^2$ is used, where r is the coefficient of correlation between estimated and actual orders (Judge et al. 1980).
15. The expectation adjustment parameter θ can only be identified if L_t and L^e_t differ. Since L^e_t always approaches L_t, a tight estimate of θ requires large variation in incoming orders from period to period. For a number of the sectors and all the retailers the variation in incoming orders is slight (recall that the retailer faces virtually constant demand). In fact, the six largest standard errors for θ are retailers. The hypothesis that expectations of customer demand are formed adaptively from past orders cannot therefore be rejected, and for one third of the sample it is supported.
16. In a study in progress, a game similar to the beer game was developed for managers of an insurance company. The game focuses on the claims-adjusting division. Like the beer game, it appears that significant underperformance comes about through misperception of the feedback structure of the system. To test the possibility of improving actual decisionmaking, the parameters of the managers' decision rules will be estimated, and the sources of poor performance fed back to the managers 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.

REFERENCES

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

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

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

Coleman, J. S., "Psychological Structure and Social Structure in Economic Models," *Journal of Business,* 59, (1986), S365-S369.

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," *Journal of Consumer Research,* 13 (1986), 25-37.

Dawes, R. M., "A Case Study of Graduate Admissions: Application of Three Principles of Human Decision Making," *American Psychologist,* 26 (1971), 180-188.

Einhorn, H. J., and R. M. Hogarth, "Behavioral Decision Theory: Processes of Judgment and Choice," *Annual Rev. Psychology,* 32 (1981), 53-88.

Einhorn, H. J., and R. M. Hogarth, "Ambiguity and Uncertainty in Probabalistic Inference," *Psychological Review,* 92 (1985), 433-461.

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

Frankel, J. A. and K. A. Froot, "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *American Economic Review,* 77 (1987), 133-153.

Grether, D., and C. Plott, "Economic Theory of Choice and the Preference Reversal Phenomenon," *American Economic Review,* 69 (1979) 623-638.

Goldberg, L. R., "Man vs. Model of Man: A Rationale, Plus Some Evidence for a Method of Improving on Clinical Inferences," *Psychological Bulletin,* 73 (1970), 422-432.

Hall, R. E., and D. W. Jorgenson, "Tax Policy and Investment Behavior," *American Economic Review,* 57, (1967) 391-414.

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

Hall, R. I., "A System Pathology of an Organization: The Rise and Fall of the Old Saturday Evening Post," *Administrative Science Quarterly,* 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.

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

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

Hogarth, R. M. and M. W. Reder, (eds), "The Behavioral Foundations of Economic Theory," *Journal of Business,* 59 (1986), S181-S505.

Holt, C., F. Modigliani, J. Muth, and H. Simon, *Planning Production, Inventories, and Workforce,* Prentice-Hall, Englewood Cliffs NJ, 1960.

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.

Kahneman, D., P. Slovic, and A. Tversky, *Judgment Under Uncertainty: Heuristics and Biases,* Cambridge University Press, Cambridge, 1982.

Leontief, W. "Theoretical Assumptions and Nonobserved Facts," *American Economic Review,* 61 (1971), 1-7.

Lopes, L. L., "Averaging Rules and Adjustment Processes: The Role of Averaging in Inference," Report 13, Wisconsin Human Information Processing Program, University of Wisconsin, Madison, 1981.

Meadows, D. L., *Dynamics of Commodity Production Cycles,* MIT Press, Cambridge MA, 1970.

Miller, G. A., E. Galanter and K. Pribram, *Plans and the Structure of Behavior,* Henry Holt, New York, 1960.

Mitchell, T. W., "Competitive Illusion as a Cause of Business Cycles," *Quarterly Journal of Economics,* 38 (1923), 631-652.

Mitchell, W. C., *Business Cycles and their Causes,* University of California Press, Berkeley, 1971.

Morecroft, J., "System Dynamics: Portraying Bounded Rationality," *Omega,* 11 (1983), 131-142.

Morecroft, J., "Rationality in the Analysis of Behavioral Simulation Models," *Management Science,* 31 (1985), 900-916.

Phelps-Brown, E. H., "The Underdevelopment of Economics," *The Economic Journal,* 82 (1972), 1-10.

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 Science,* 24 (1978), 827-835.

Richardson, G. P., *The Evolution of the Feedback Concept in American Social Science,* Unpublished Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge MA, 1984.

Samuelson, P. A., "Interactions Between the Multiplier Analysis and the Principle of Acceleration," *The Review of Economic Statistics,* 21 (1939), 75-78.

Simon, H. A., "Rational Decisionmaking in Business Organizations," *American Economic Review,* 69 (1979), 493-513.

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

Simon., H. A., "The Behavioral and Rational Foundations of Economic Dynamics," *Journal of Economic Behavior and Organization,* 5, (1984), 35-55.

Slovic, P. and S. Lichtenstein, "Preference Reversals: A Broader Perspective," *American Economic Review,* 73 (1983), 596-605.

Smith, V. L., "Experimental Methods in the Political Economy of Exchange" *Science,* 234, (10 October 1986), 167-173.

Smith, V. L., "Microeconomic Systems as an Experimental Science," *American Economic Review,* 72 (1982), 923-955.

Sterman, J. D., "A Behavioral Model of the Economic Long Wave," *Journal of Economic Behavior and Organization,* 6 (1985), 17-53.

Sterman, J. D., "Testing Behavioral Simulation Models by Direct Experiment," *Management Science,* 33 (December) 1987a.

Sterman, J. D., "Misperceptions of Feedback in Dynamic Decisionmaking," *Organizational Behavior and Human Decision Processes.* Forthcoming. Also D-3876-2, System Dynamics Group, Sloan School of Management, MIT, Cambridge MA, 1987b.

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

Tversky, A. and D. Kahneman, "Judgment Under Uncertainty: Heuristics and Biases," *Science,* 185 (27 September 1974), 1124-1131.

Tversky, A. and D. Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science,* 211 (30 January 1981), 453-458.

Tversky, A. and D. Kahneman, "Rational Choice and the Framing of Decisions," *Journal of Business,* 59 (1986), S251-S278.

Zarnowitz, V., *Orders, Production, and Investment – a Cyclical and Structural Analysis,* National Bureau of Economic Research, New York, 1973.

Zeckhauser, R., "Behavioral versus Rational Economics: What You See is What You Conquer," *Journal of Business,* 59 (1986), S435-S449.

Figure 1. The generic stock-management system.

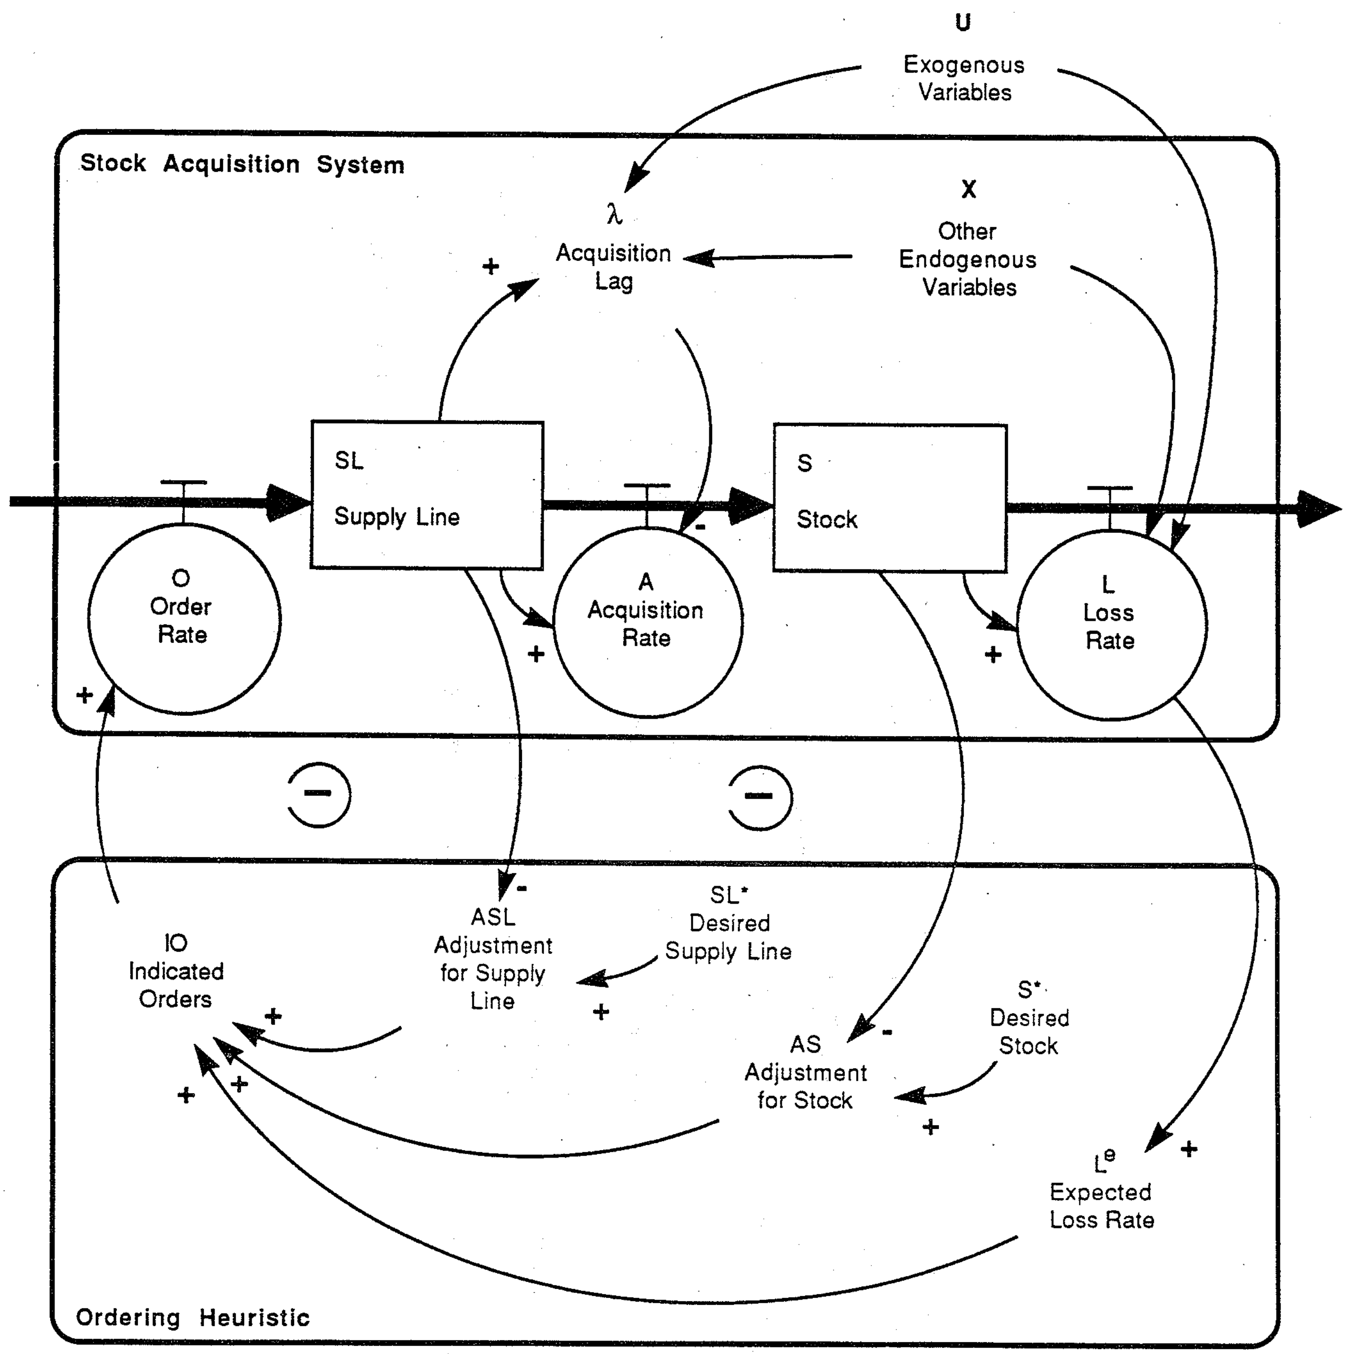

Table 1. Examples of stock-management systems

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	Vacancies & 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

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 shiping 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.* Orders 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.

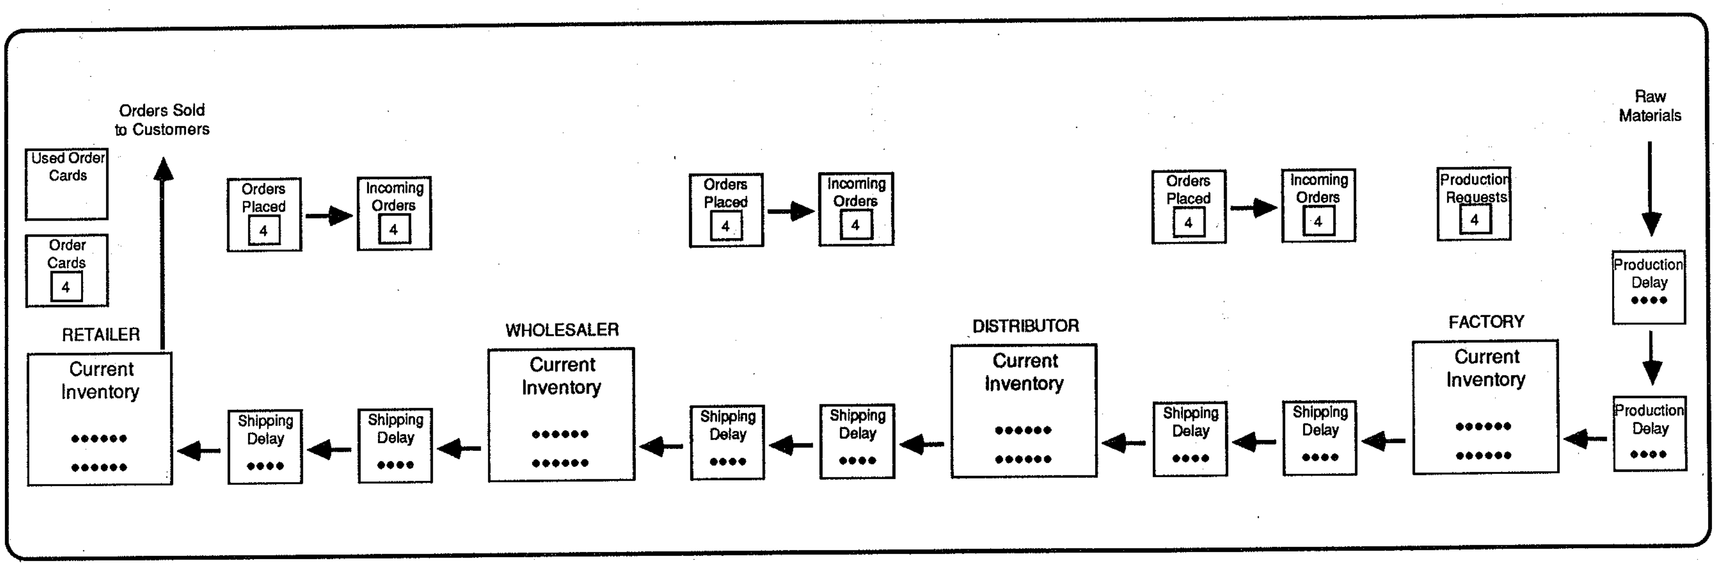

Figure 3. Customer Orders. In week 5 customer orders rise from 4 to 8 cases per week. Compare against the oscillations in the subjects' orders (figure 4).

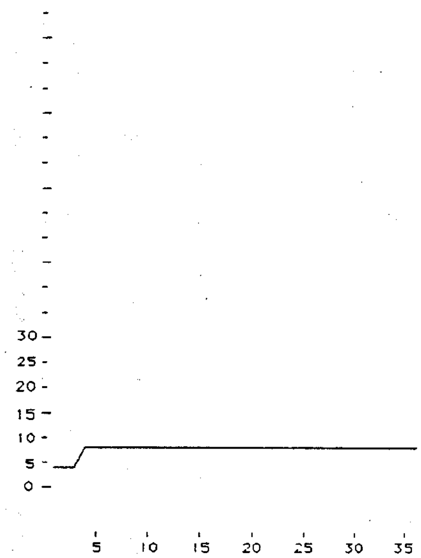

Figure 4a. Key to experimental results (figure 4b).

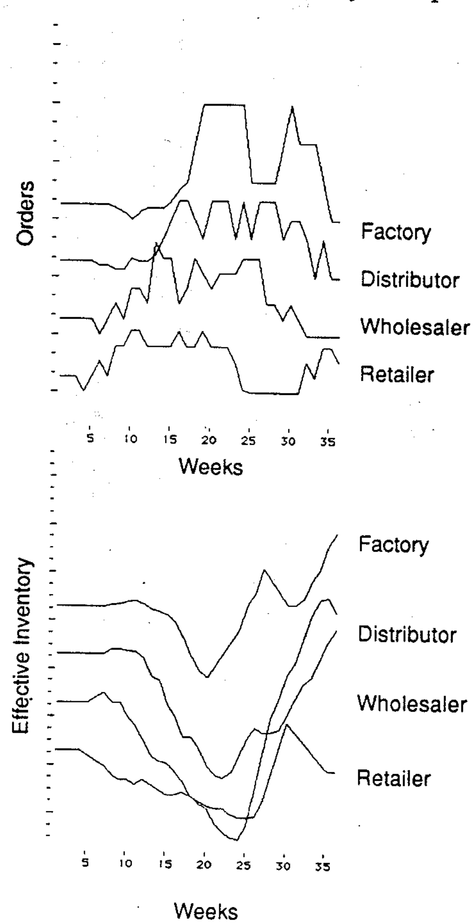

Orders placed by sector.
From bottom to top,
R, W, D, F, each offset by 15 cases/week.
Major tick-marks=15 cases/week.
Minor tick-marks=5 cases/week.
Initial orders=4 cases/week in all sectors.

Effective Inventory by sector.
Effective Inventory=Inventory-Backlog.
From bottom to top,
R, W, D, F, each offset by 40 cases.
Major tick-marks=40 cases.
Minor tick-marks=10 cases.
Initial inventory=12 cases in all sectors.

Figure 4b. Typical Experimental Results

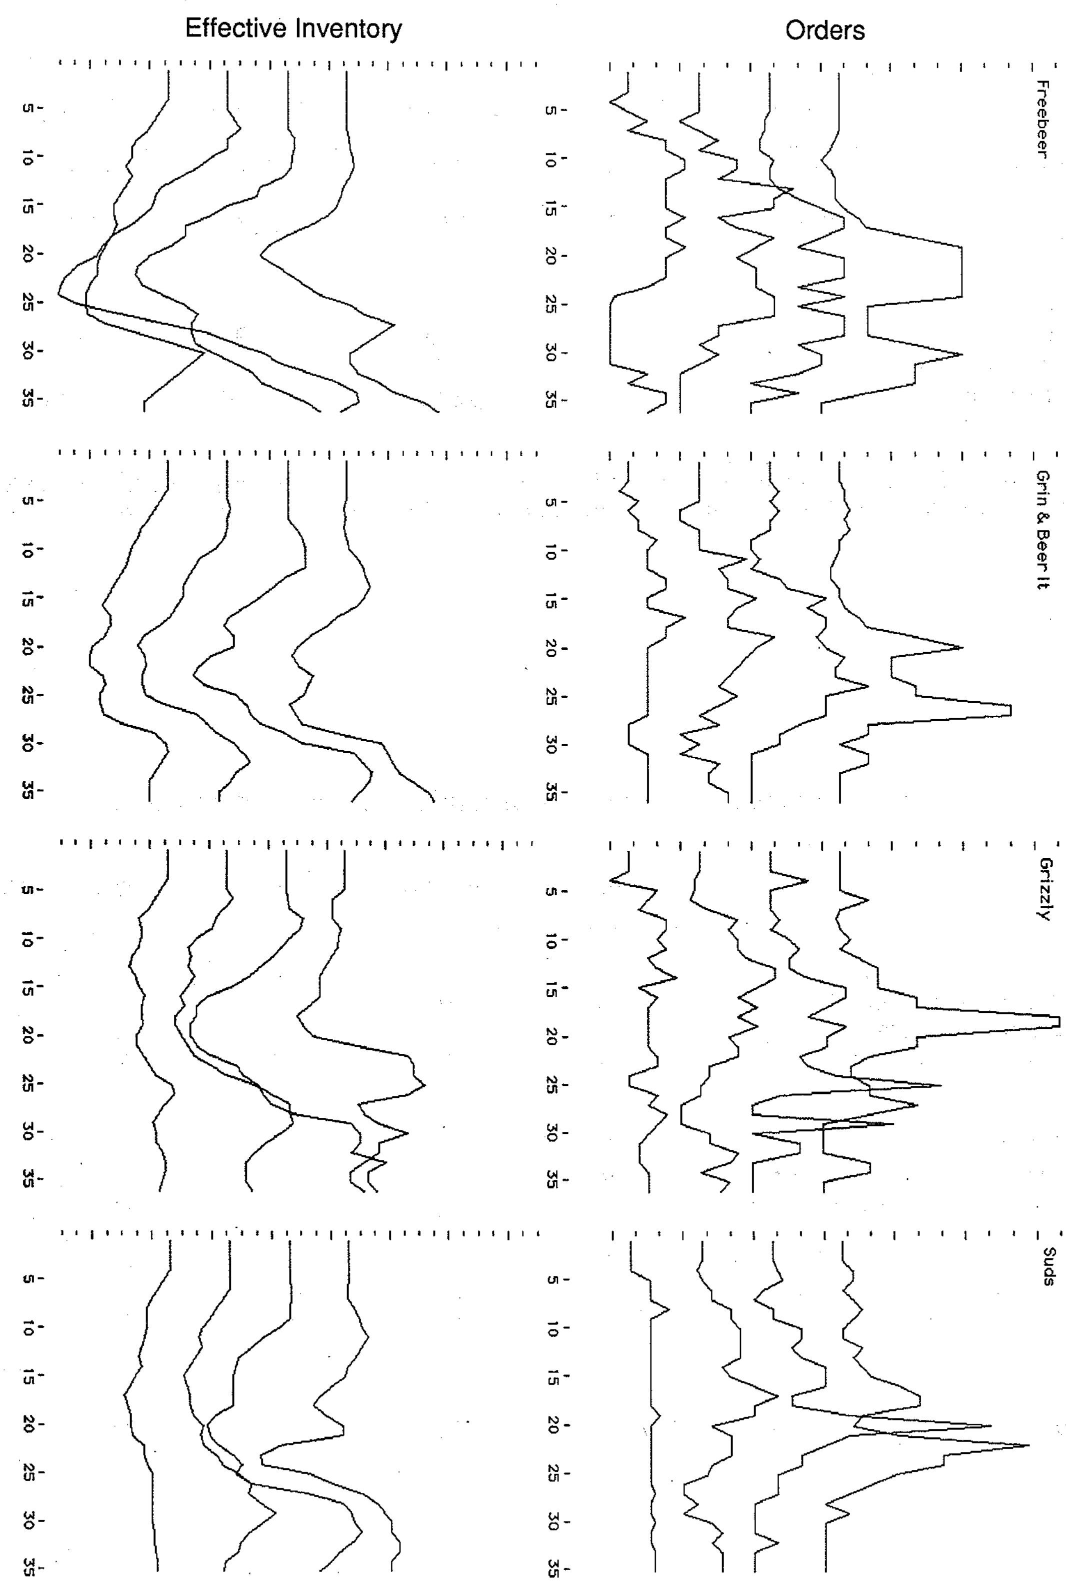

Table 2. Comparison of experimental and benchmark costs.

	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	$p<.000+$	$p<.001$	$p<.000+$	$p<.000+$	$p<.000+$

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

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

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	0.85
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: .005; b: .01; c: .025 level (1-tailed t-test [since parameters must be ≥ 0])

Figure 5. Typical sample of subjects' post-play judgments of customer orders. Compare against actual customer orders (Figure 3).

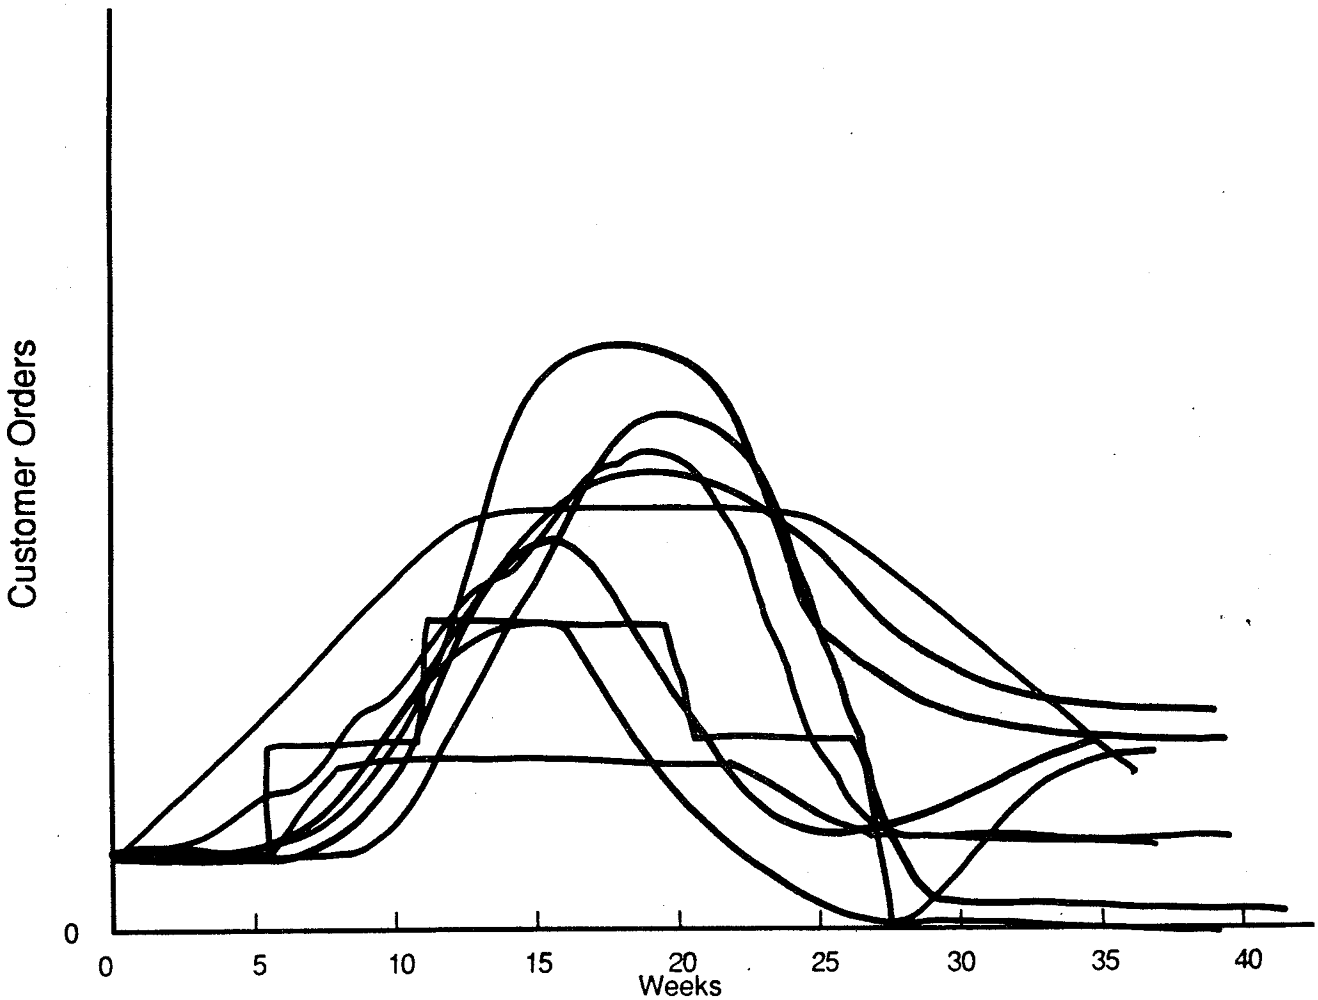

Figure 6. Three stages of industrial production. Ratio to trend, 1947-1987. Note the growing oscillation, amplification, and phase lag from consumer goods to intermediate goods to materials production. Source: See Table 5.

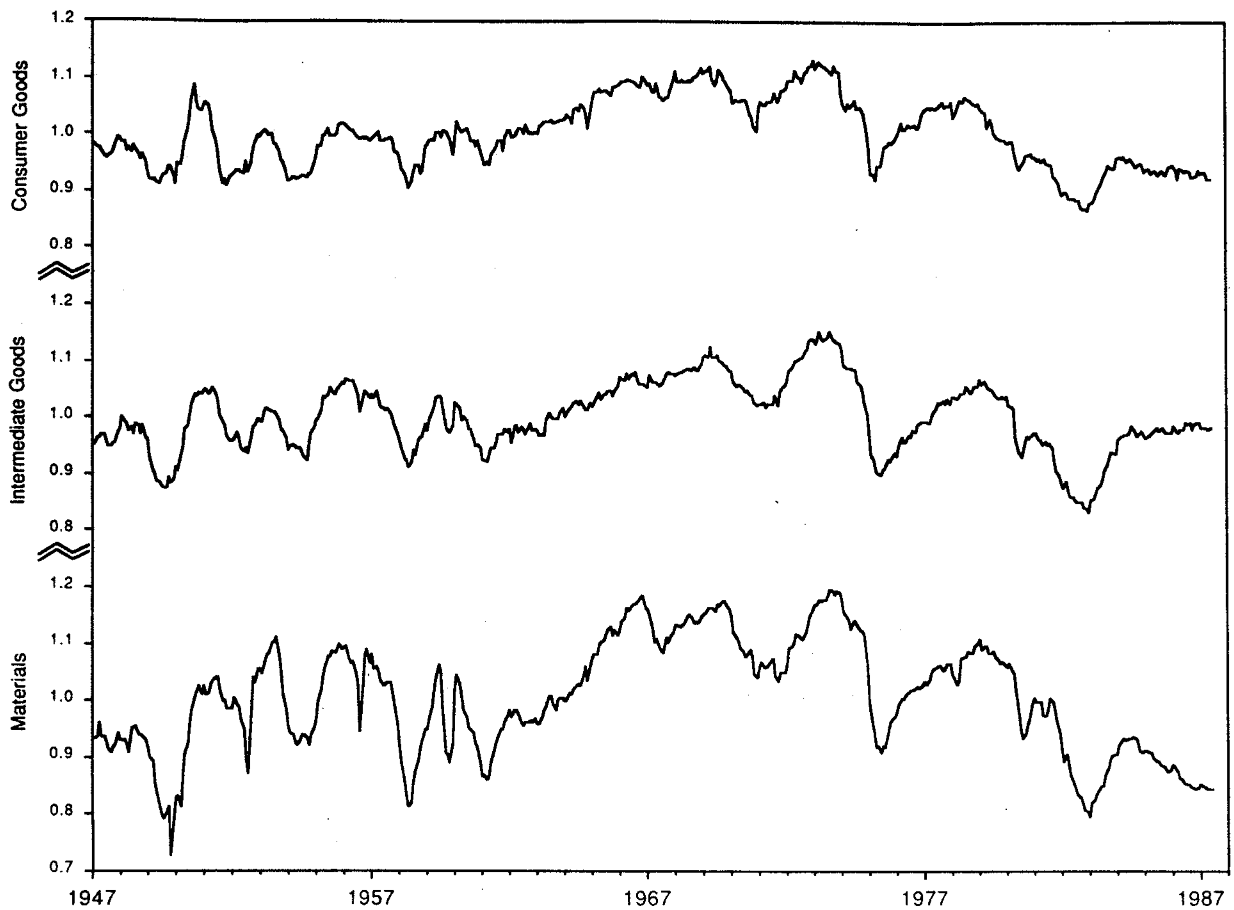

Table 5. Amplification and phase lag in three stages of production.

Standard Deviation (%)

Consumer Goods	6.22
Intermediate Goods	6.31
Materials	10.00

Phase Lag (months)

Materials – Intermediate Goods	1.5
Intermediate Goods – Consumer Goods	2.1

Source: Federal Reserve Board, industrial production index for Consumer Goods, Intermediate Goods, and Materials; monthly data, 1947.1-1987.5. Figure 6 and table above show detrended data. The ratio to trend $R_t = P_t/T_t$; P_t=Production and T_t=Trend. Trend $T_t = \exp(a+bt)$ where a, b are determined for each series by linear regression on the log of the production index: $\ln(P_t) = a + b^*t$.