HD28

.M414

no. 1945-87

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

DETERMINISTIC CHAOS IN MODELS OF HUMAN BEHAVIOR:
METHODOLOGICAL ISSUES AND EXPERIMENTAL RESULTS

JOHN D. STERMAN

OCTOBER 1987

WP 1945-87

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

DETERMINISTIC CHAOS IN MODELS OF HUMAN BEHAVIOR:
METHODOLOGICAL ISSUES AND EXPERIMENTAL RESULTS

JOHN D. STERMAN

OCTOBER 1987

WP 1945-87

**Deterministic Chaos in Models of Human Behavior:
Methodological Issues and Experimental Results**

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

Forthcoming in the
Special Issue on Instabilities and Chaos
System Dynamics Review

October 1987

ABSTRACT

Recent work has shown that several well-known models in the system dynamics literature contain previously unsuspected regimes of deterministic chaos. Two of the most extensively analyzed are Sterman's model of the economic long wave and the production-distribution model of the Beer Distribution Game. The significance of these theoretical developments hinges on whether the chaotic regimes lie in the realistic region of parameter space. Further there are major questions regarding the descriptive accuracy of models of human systems which exhibit chaos. Data limitations and the inability to conduct controlled experiments mean empirical studies at the aggregate level are not likely to resolve these questions. An alternative approach is based on laboratory experiments in which models provide a simulated environment for the study of human decisionmaking behavior. Recently laboratory experiments have been conducted to analyze decisionmaking behavior in the long wave model and the Beer Distribution Game. This paper describes the experiments and shows that the behavior of the subjects is explained well with a simple heuristic long used in system dynamics modeling and well grounded in behavioral decision theory. The parameters of the proposed decision rule are estimated econometrically for each subject. The parameters which characterize a significant minority of the subjects are shown to produce chaos. This direct experimental evidence that chaos can be produced by the decisionmaking behavior of real people has important implications for the formulation, analysis, and testing of models of human behavior.

Recent work has shown that several well-known models in the system dynamics literature contain previously unsuspected regimes of deterministic chaos. The work of Day (1982a, 1982b) provides an early example of chaotic behavior in economic models, while Mosekilde and others (1987) have developed corporate models which exhibit chaos. Two of the most extensively analyzed such models are Sterman's model of the economic long wave (Sterman 1985, 1986, Rasmussen, Mosekilde, and Sterman 1985) and the production-distribution system or Beer Distribution Game (Forrester 1961, Jarmain 1963, Sterman 1984, Mosekilde and Larsen, this issue). While the demonstration that chaos can be endogenously produced in these systems is an important theoretical development, the significance of the results hinges in large measure on whether the chaotic regimes lie in the realistic region of parameter space or whether they are mathematical curiosities never observed in the real system. Further there are major questions regarding the descriptive accuracy of the decision rules postulated in the models of human systems developed to date which contain strange attractors. The practical significance of chaos and other phenomena such as self-organization in policy-oriented modeling remains unclear until it can be determined that these phenomena can occur in models whose decision rules are grounded in empirical study of the actual decision processes of the agents. It is difficult if not impossible to resolve such issues by appeal to the aggregate empirical data. In the case of the long wave, for example, there have been at most five long-wave cycles since the industrial revolution, too few for statistically reliable results. Worse, the data required are simply unavailable, and much of it is corrupted by measurement error (Chen, this issue).

An alternative approach is based on laboratory experiments in which models provide a simulated environment for the study of human decisionmaking (Sterman 1987a). Recently, the long wave model and Beer Distribution Game (traditionally used as a teaching rather than as a research tool) have been analyzed as experiments in dynamic decisionmaking

(Sterman 1987b, 1987c). These experiments were conducted primarily to study the heuristics with which people manage a complex dynamic environment. In both experiments the decision task of the subject was to manage a stock in the face of losses, delays in acquiring new units, multiple feedbacks and other environmental disturbances.

This paper describes the experiments and shows that decisionmaking behavior in the two experiments is significantly suboptimal. The departures from optimal behavior are systematic, suggesting subjects use common heuristics for stock management. The experiments show that the behavior of the subjects can be modeled well with a simple stock-management heuristic long used in system dynamics modeling and well grounded in behavioral decision theory. The parameters of the proposed heuristic are estimated econometrically for each subject. The systems are then simulated with the estimated parameters. The parameters which characterize a significant minority of the subjects are shown to produce chaos. This direct experimental evidence that chaos can be produced by the decisionmaking behavior of real people has important implications for the formulation, analysis, and testing of models of human behavior.

The Stock Management Problem

The regulation of a stock or system state is one of the most common dynamic decisionmaking tasks. The manager seeks to maintain a quantity at a target level, or at least within an acceptable range. Stocks cannot be controlled directly but rather are influenced by altering the stock's inflow and outflow rates. Typically a manager must set the inflow rate so as to compensate for losses from the stock and to counteract environmental disturbances which may 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 and compensating for variations in credit demand.

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 :

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

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

$$L_t = f_L(S_t, X_t, U_t). \quad (2)$$

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:

$$A_t = f_A(SL_t, \lambda_t). \quad (3)$$

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

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

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 numerous exogenous forces including noise and nonstationarity of the underlying equilibrium.

Consistent with the behavioral foundations of system dynamics modeling (Morecroft 1983, 1985, Sterman 1987a), behavioral decision theory (Hogarth 1987, Tversky and Kahneman 1974), and the theory of bounded rationality (Simon 1979, Cyert and March 1963), the proposed decision rule 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 which any stock management heuristic must include:

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 (e.g. shipments from an inventory of finished goods) or decay (e.g. the depreciation of plant and equipment). Failure to replace losses would cause the stock to fall below the desired level, creating steady-state error.

2. *Stock adjustment.* Errors in forecasting losses or changes in the desired stock demand 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 rules 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. Failure to account for the supply line results in overcorrection and instability. Consider cooking dinner on an electric range. If one turns the heat down just as the pot comes to a boil the supply line of heat in the coils of the range will continue to heat the pot, boiling it over and ruining dinner.

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

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

where IO is the indicated order rate, the rate indicated by other pressures.¹

The indicated order rate is based on the anchoring and adjustment heuristic (Tversky and Kahneman 1974). Anchoring and adjustment is a common judgmental 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. 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, Lopes 1981, Hines 1987, Sterman 1987d). Here the anchor is the expected loss rate L^e . Adjustments are then made in response to discrepancies between the desired and actual stock AS and between the desired and actual supply line ASL:

$$IO_t = L^e_t + AS_t + ASL_t. \quad (6)$$

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 $L^e_t = \gamma L_{t-1} + (1-\gamma)L^*$, $0 \leq \gamma \leq 1$, adaptive expectations $L^e_t = \theta L_{t-1} + (1-\theta)L^e_{t-1}$, $0 \leq \theta \leq 1$, and extrapolative expectations, $\Delta L^e_t = \sum \omega_i \Delta L_{t-i}$, 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:

$$AS_t = \alpha_S (S^*_t - S_t) \quad (7)$$

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

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

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:

$$SL^*_t = \lambda^e_t * \Phi^*_t. \quad (9)$$

The adjustment for the supply line creates a second negative feedback loop which avoids overordering and also compensates for changes in the acquisition lag λ . If λ rises, for example, ASL induces sufficient additional orders to restore $A = \Phi^*$. There are a variety of possible representations for λ^e and Φ^* , ranging from constants through sophisticated forecasts.

Experiment I: The Economic Long Wave

The experimental protocol for the long wave model is described in Sterman and Meadows 1985 and Sterman 1987a, 1987b). The methodological foundations of experimental economics are discussed in the seminal work of Smith (1982) and Plott (1986). The model represents the aggregate capital-producing sector of the economy (figure 2). The dynamic

hypothesis behind the long wave can be stated in two parts. First, due to construction lags, the process by which individual firms adjust production capacity to demand is inherently oscillatory. In isolation an individual firm would produce damped fluctuations with a period of about 15-25 years in response to a shock. Second, in the aggregate the capital-producing sector of the economy orders and acquires capital from itself. This multiplier effect or 'self-ordering' creates a positive feedback loop which destabilizes the oscillatory tendencies of individual firms, changing the damped oscillation to a limit cycle with a 40-60 year period. Simulation and formal analysis confirm the dynamic hypothesis (Sterman 1985; Rasmussen, Mosekilde and Sterman 1985). As the self-ordering loop becomes stronger, damping drops rapidly. The system goes through a Hopf bifurcation and produces a limit cycle. Further increases in the strength of self-ordering proceed through period doublings and ultimately to chaos.

In the experiment the model is transformed into a game in which subjects manage the capital-producing sector of a simple economy. Subjects must balance the supply of and demand for capital by adjusting their production capacity to meet desired production. Orders for capital arise from the exogenous consumer goods sector and from the subject's own orders. The equations underlying the game reflect a simplified form of the original model but preserve its essential structure and dynamics. In the game decisions are made in discrete time intervals representing two years, and the model becomes a third-order difference equation system.

The equations of the game are given below. Production PR is the lesser of desired production DP or Production Capacity PC. Capacity is proportional to the capital stock. The capital/output ratio κ is one period (two years):

$$PR_t = \text{MIN}(DP_t, PC_t) \quad (10)$$

$$PC_t = K_t / \kappa. \quad (11)$$

The capital stock of the capital sector is augmented by acquisitions AK and diminished by depreciation CD . Depreciation is proportional to the stock. The average lifetime of the capital stock τ is 10 periods (20 years):

$$K_{t+1} = K_t + (AK_t - CD_t) \quad (12)$$

$$CD_t = K_t/\tau. \quad (13)$$

The acquisition of capital by both the capital and goods sectors (AK and AG) depends on the supply line of unfilled orders each has accumulated (the backlogs BK and BG) and the fraction of demand satisfied FDS . Each period both the capital and goods sectors acquire the full supply line of unfilled orders unless the capital sector is unable to produce the required amount. In this case acquisition by each sector is reduced in proportion to the shortfall. The acquisition lag λ is thus $1/FDS$, and $AK+AG=PR$ at all times, ensuring that output is conserved:

$$AK_t = BK_t * FDS_t \quad (14)$$

$$AG_t = BG_t * FDS_t \quad (15)$$

$$FDS_t = PR_t/DP_t \quad (16)$$

$$DP_t = BG_t + BK_t. \quad (17)$$

The supply lines of unfilled orders for each sector BG and BK are augmented by orders for capital placed by each sector and emptied when those orders are delivered:

$$BK_{t+1} = BK_t + (OK_t - AK_t) \quad (18)$$

$$BG_{t+1} = BG_t + (OG_t - AG_t). \quad (19)$$

New orders placed by the goods sector are the exogenous input to the system to which the subject of the experiment must respond by choosing an appropriate amount of capital to order for their own use:

$$OG_t = \text{exogenous} \quad (20)$$

$$OK_t = \text{determined by subject.} \quad (21)$$

Subjects are responsible for only one decision – how much capital to order. The goal

of the subjects in making these decisions is to minimize their total score for the trial. The score is defined as the average absolute deviation between desired production DP and production capacity PC over the T periods of the experiment:

$$S = (1/T) \sum_{t=0}^T |DP_t - PC_t|. \quad (22)$$

The score indicates how well subjects balance demand and supply. Subjects are penalized equally for both excess demand and excess supply.

The experiment is implemented on IBM PC-type microcomputers. A 'game board' is displayed on the screen providing subjects with perfect information. Color graphics and animation highlight the flows of orders, production, and shipments to increase the transparency of the structure. The conversion of the original model to a form suitable for experimental testing is described in Sterman 1987a.

The subject population (N=49) consisted of MIT undergraduate, master's and doctoral students in management and engineering, many with extensive exposure to economics and control theory; scientists and economists from various institutions in the US, Europe, and the Soviet Union; and business executives experienced in capital investment decisions including several company presidents and CEOs.

Typical experimental results are shown in figure 3. All trials begin in equilibrium. In period 3 there is a one-time step input of 10% in the orders of the goods sector. The optimal response returns the system to equilibrium within 6 periods, producing a score of 19. In contrast the vast majority of subjects produced significant oscillations (table 1). The average score for the sample was 591; the lowest was 77.

Next the proposed stock management heuristic (equations 5-9) was tested against the ordering behavior of the subjects. Adapting the heuristic to the experimental context is straightforward. The stock to be managed is the capital stock K. The desired stock is

proportional to desired production. The loss rate L is simply the depreciation of the capital stock CD . The desired supply line DSL was specified according to equation 9, with the expected acquisition lag $\lambda^e = \lambda = 1/FDS$ and desired throughput $\Phi^* = CD$. Allowing an additive disturbance term ϵ , the proposed ordering rule for capital investment becomes:

$$OK_t = \text{MAX}(0, CD_t + AS_t + ASL_t + \epsilon_t) \quad (23)$$

$$AS_t = \alpha_S(DK_t - K_t) \quad (24)$$

$$ASL_t = \alpha_{SL}(DSL_t - BG_t) \quad (25)$$

$$DSL_t = \lambda^e * CD_t = (1/FDS_t) * CD_t. \quad (26)$$

To test the rule only the two adjustment parameters α_S and α_{SL} need be estimated. All other data required to determine orders are presented directly to the subjects. Maximum likelihood estimates of the parameters for each trial were found by grid search of the parameter space, subject to the constraints $\alpha_S, \alpha_{SL} \geq 0$. Assuming the errors ϵ are independent, identical, and normally distributed then the maximum likelihood estimates of such nonlinear functions are given by the parameters which minimize the sum of squared errors. Such estimates are consistent and asymptotically efficient, and the usual measures of significance such as the t -test are asymptotically valid.

Estimates for 49 trials together with t -statistics are given in table 2. The model's ability to explain the ordering decisions of the subjects is excellent. R^2 varies between 33% and 99+%, with an overall R^2 for the pooled sample of 85%.² All but two of the estimated capital stock adjustment parameters are highly significant. The supply line adjustment parameter is significant in 22 trials.

Sterman 1987b analyzes the estimation results and identifies several 'misperceptions of feedback' which are responsible for the subjects' poor performance. One of these is the tendency for subjects to give insufficient attention to the supply line, causing subjects to

continue ordering even after the construction pipeline contains sufficient units to correct any stock discrepancy. The present concern, however, is the relationship between the estimated parameters and the regimes of behavior in the model. The estimated parameters characterize the decision making behavior of actual people. When simulated in the model are the estimated decision rules of the subjects inherently stable, or do they produce limit cycles, period multiples, or chaos?

Table 2 also indicates the mode of behavior produced by simulation of the decision rule with the estimated parameters. The parameters estimated for thirty subjects (61%) are stable. Most of these produce overdamped behavior of the capital stock in response to the step input. Seven parameter sets produce limit cycles of period 1, and 2 produce period multiples. The parameters which characterize ten subjects (20%) produce chaos. Inspection of table 2 shows that the subjects whose parameters are stable performed best in the task while those whose parameters produce limit cycles and chaos generally had the highest scores.

Figure 4 shows the phase portrait for several of the more interesting parameter sets. Note the period 5 cycle produced by $(\alpha_S, \alpha_{SL}) = (.46, .33)$. The apparent period 2 cycle produced by $(.42, 0)$ is, upon closer inspection, seen to be chaotic.

Figure 5 locates the modes of behavior in parameter space. Consistent with both intuition and formal stability analysis, stability is enhanced by small values of α_S and large values of α_{SL} . More aggressive attempts to correct the discrepancy between the desired and actual capital stock are destabilizing: by ordering more aggressively the self-ordering loop causes a larger increase in total demand, thus exacerbating disequilibrium and encouraging still larger orders in future periods. Conversely, more aggressive response to the supply line is stabilizing by constraining orders as the supply line fills.

Thus the periodic and chaotic attractors tend to occur for large values of α_S and small values of α_{SL} . Yet the route to chaos is apparently more complex. Consider the route to chaos as α_S increases for $\alpha_{SL}=0$. Stability gives way to period one limit cycles, then to period 2 at (.42, 0), and ultimately to chaos for $\alpha_S \geq .55$. Yet the parameters which produce the slightly chaotic period 2 are surrounded on the ray $\alpha_{SL} = 0$ by parameters which produce period 1 limit cycles. Similarly, the period 5 attractor produced by the parameters (.46, .33) lies between the region of stability and the region of period 1 behavior.

The results show that the parameters which characterize actual decisionmaking performance in a common and important macroeconomic situation can produce a wide range of behavior, including chaos.

Experiment II: The Production-Distribution System

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 6). 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 cumulative team costs over the length of the game. Inventory holding costs for each sector 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).

The protocol for the experiment is described in Sterman 1987c. The game is initialized in equilibrium. Each inventory contains 12 cases. 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.

The results reported here were drawn from four dozen games (192 subjects) collected over a period of four years. A computer model of the game was used to test the records for consistency. Trials in which there were accounting 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, providing 44 subjects.³ 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.

Typical experimental results are shown in figure 7 and summarized in table 3. The results show the behavior of the subjects to be far from optimal. The average team cost is ten times greater than optimal. Like the macroeconomic experiment, the results exhibit several strong 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.

2. *Amplification*: The amplitude and variance of 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. 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.

3. *Phase lag*: The peak order rate tends to occur later as one moves from the retailer to the factory. 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.

Next 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. 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 1987d, Frankel and Froot 1987), and they are one of the simplest formulations for expectations flexible enough to adapt to a nonstationary process. Because the subjects lack the information and time to determine optimal inventory or supply line levels the desired stock S^* and desired supply line SL^* are both assumed to be constants which must be estimated.

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

$$O_t = \text{MAX}(0, IO_t) \quad (27)$$

$$IO_t = L^e_t + AS_t + ASL_t \quad (28)$$

$$L^e_t = \theta L_{t-1} + (1-\theta)L^e_{t-1}, \quad 0 \leq \theta \leq 1 \quad (29)$$

$$AS_t = \alpha_S(S^* - S_t) \quad (30)$$

$$ASL_t = \alpha_{SL}(SL^* - SL_t). \quad (31)$$

Defining $\beta = \alpha_{SL}/\alpha_S$ and $S' = S^* + \beta SL^*$ yields

$$IO_t = L^e_t + \alpha_S(S' - S_t - \beta SL_t). \quad (32)$$

Note that since S^* , SL^* , α_{SL} and α_S are all ≥ 0 , $S' \geq 0$. 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 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 by the subjects. If $\beta = 1$, then the subjects fully recognize the supply line and do not double order. If $\beta = 0$, then orders placed are forgotten until they arrive, encouraging overordering and instability.⁴

In contrast to the macroeconomic experiment, the stock management decision rule as

implemented in the beer game contains four parameters to be estimated (θ , α_S , S' , and β). The parameters were estimated by the same procedure used in the long wave experiment. An additive disturbance term is assumed, and the disturbances are assumed to be i.i.d. normal. Estimates for each sector of each trial were found by grid search of the parameter space subject to the constraints $0 \leq \theta \leq 1$ and $\alpha_S, S', \beta \geq 0$.⁵

The estimated parameters are displayed in table 4. The mean R^2 is 71%; R^2 is less than 50% for only 6 of 44 subjects.⁶ A large majority of the estimated parameters are significant. As in the long wave experiment the decision rule is an excellent model of the decisionmaking behavior of the subjects. Sterman 1987c relates the estimated parameters to the dysfunctional performance in the task. It is noteworthy that the same misperceptions of feedback structure which cause oscillation in the macroeconomic experiment are responsible for the poor performance of the subjects in the beer game. Specifically, the supply line of unfilled orders is given far too little weight by most subjects.

Next the beer game was simulated using the stock adjustment rule with the estimated parameters. Each trial involves four subjects, each of whom could have different parameters. In the simulations, however, the same parameters were assumed for each sector, vastly simplifying the analysis.⁷

Table 4 indicates the mode of behavior produced by each set of parameters. Thirty-one parameter sets (70%) produce stable behavior. One each produces limit cycles of period 1, period 5, and period 12. Ten yield chaos. Since the beer game is a 19th order system the behavior produced by the system is exceedingly complex and there are many possible projections of phase space. Figure 8 shows several of the more interesting phase portraits. Note, for example, the different shape of the attractor produced by simulation with the

parameters of the Bassbeer Factory ($\theta, \alpha_s, \beta, S' = 1, .65, .40, 15$) compared to those of the Heineken3 Factory ($\theta, \alpha_s, \beta, S' = 0, .3, .15, 17$). The parameters estimated for the Coors Factory ($\theta, \alpha_s, \beta, S' = .25, .3, 0, 18$) reveal a cycle of period 12 when distributor inventory is plotted against factory inventory. The retail sector of the same team, however, produces a limit cycle of period 1. The difference in mode is probably due to the integer constraint on orders.

Figure 9 shows the distribution of the modes in parameter space. The relationships between the adjustment parameters and stability are similar to those apparent in the macroeconomic experiment, and for similar reasons. In general, larger values of α_s and smaller values of β are destabilizing. That is, to the extent the subjects ignore the supply line of unfilled orders ($\beta \approx 0$) a stock discrepancy will cause orders to be placed each period even after sufficient orders are in the supply line, leading to excess inventories and oscillation. Such overordering is exacerbated by more aggressive stock adjustment ($\alpha_s \gg 0$) since more will be ordered in response to a given discrepancy between the desired and actual stock. Figure 9 also indicates that in general the smaller the desired inventory S' the less stable the system. Again, intuition readily explains why this is so. Consider the response of the retailer to an unanticipated increase in customer demand. Retail inventory falls due to the lags in receiving beer. The retailer will place additional orders with the wholesaler through the stock adjustment term. If the wholesaler has adequate inventory to fill incoming orders ($S' \gg 0$) incoming orders can all be filled and the retailer's acquisition lag remains constant. If, however, the wholesaler runs out of inventory the retailer finds orders placed fail to arrive. The delay in receiving beer lengthens, the retailer's stock drops still further, and still more orders swell the supply line in a vain attempt to replenish inventory. When sufficient beer finally does arrive, the retailer finds inventory overshooting the desired level by huge margins

(figure 7). S' plays such a crucial role in determining the stability of the system because it determines how close the system operates to the fundamental nonlinear constraint that shipments can only be made if inventory exists. Small values of S' make it more likely the system will run out of inventory thus entering the unstable region in which the acquisition lag λ is long and variable. Those subjects who maintained larger target inventory levels were thus more stable on average than those who attempted to cut costs by reducing their buffer stock.

There appears to be no simple relationship between the parameter θ and the modes of the system. θ controls the speed with which the adaptive forecast of incoming orders is adjusted and has two opposing effects on stability. Faster adjustment of the demand forecast ($\theta \approx 1$) increases the gain between orders placed and incoming orders. Shocks are transmitted up the distribution chain more readily, destabilizing the system by causing larger inventory excursions upstream in response to a given change in orders downstream. On the other hand, sluggish adjustment ($\theta \approx 0$) means variations in incoming orders are not quickly countered by the replacement term of the ordering rule, causing inventory to decline further than it would if replacement of incoming orders was swift. As a result, inventories may be exhausted, moving the system into the less stable zone in which backlogs build up and the acquisition lag lengthens.

As in the macroeconomic experiment the results show an unexpectedly rich array of behavior is generated by the interaction of the decision processes of real people with the feedback structure of the environment.

Discussion

The discovery of nonlinear phenomena such as deterministic chaos in the physical world naturally motivates the search for similar behavior in the world of human behavior. Yet the social scientist faces difficulties in that search which do not plague the physicist, at least

not to the same degree. Aggregate data sufficient for strong empirical tests simply do not exist for many of the most important social systems. Social systems are not easily isolated from the environment. The huge temporal and spatial scales of these systems, vast number of individual actors, considerations of cost and ethical concerns make controlled experiments on the systems themselves difficult at best. Finally, the laws of human behavior are not as stable as the laws of physics.

The social scientist is thus left with two main options in the quest to understand the relevance of nonlinear phenomena in the domain of human behavior. Formal models of social systems may be constructed and analyzed. Previous work of this type has demonstrated the existence of chaotic modes in a wide array of social and economic models. However, analysis of these models frequently showed that the chaotic regime lay far outside the plausible region of parameter space. Worse, without empirical tests of the decision rules the verisimilitude of the models was open to doubt. In the absence of empirical data the relevance of nonlinear phenomena remained questionable.

The second approach is to develop laboratory experiments with simulated social systems. Since experiments on actual firms and national economies are infeasible, simulation models of these systems must be used to explore the decisionmaking heuristics of real people. Such experiments create 'microworlds' in which the subjects face physical and institutional structures, information, and incentives which mimic (albeit in a simplified fashion) those of the real world.

The experiments described here presented subjects with a straightforward stock-adjustment task. Results show that the subjects' behavior can be modeled with a high degree of accuracy by a decision rule long used in system dynamics models. The decision rule is consistent with a vast body of empirical knowledge developed in behavioral decision theory, organizational studies, and psychology. The stock-adjustment heuristic appears to

be robust: the feedback structure, information available, loss function, time available, "cover story" and many other dimensions of the experimental environment differed markedly. Performance in both settings was decidedly suboptimal. The same misperceptions of the feedback structure of the system are apparent in both cases.

The experiments reported here demonstrate that chaos can in fact be produced by the decisionmaking processes of real people. The results are significant in several respects. The demonstration that chaos can be produced by the decisionmaking heuristics people actually use strengthens the argument for the universality of these phenomena. Chaos thus appears to be a common mode of behavior not only in physical systems but in social and economic systems as well. These results do *not* show that, say, the evolution of GNP in the United States has in fact been chaotic. Further empirical work is required before such questions can be settled. But the results suggest modelers can ignore nonlinear dynamics only at their peril. Models of economic and social dynamics should portray the processes by which disequilibrium conditions are created and dissipated, not assume that the economy is in or near equilibrium at all times or that adjustment processes are stable. Models should be formulated so that they are robust in extreme conditions, since it is the nonlinearities necessarily introduced by robust formulations that crucially determine the modes of the system. These principles have long been central to system dynamics (Forrester 1961, Forrester and Senge 1980, Richardson and Pugh 1981) and behavioral economics (Day 1984) and are important even for systems which do not contain limit cycles or strange attractors (Sterman 1985b).

At the same time a number of questions regarding the practical significance of chaos and other nonlinear phenomena in social system modeling remain unresolved. Real social systems are bombarded by broadband noise, and it is well known that such random shocks severely degrade the point-predictability of most systems (Forrester 1961, App. K). Does

the existence of stochastic shocks swamp the uncertainty in trajectories caused by chaos? How does the existence of chaotic regimes in a model influence its response to policies, and the predictability of that response? Particularly troubling here is the apparent nonuniformity of the distribution of modes in parameter space. Mosekilde and Larsen's (this issue) carpet bombing experiments with the simulated beer game show that there are islands of chaos in parameter space, and that the boundary between the periodic and chaotic solutions is fractal, a finding consistent with analysis of such classic systems as Duffing's equation (Thompson and Stewart 1986). The response of stability to parameter variations may not be monotonic or even predictable in such systems, raising major questions for policy analysis. The development of principles for policy design in such systems is a major area for future research.

Similarly, chaos is a steady-state phenomenon which manifests over very long time frames, but many policy-oriented models are concerned with transient dynamics and nearly all with time horizons much shorter than those used in the analysis of chaotic dynamics (e.g. Mosekilde et al. 1986, Weidlich et al. this issue). Over such extended time horizons the parameters of the system cannot be considered static but will themselves evolve with learning and evolutionary pressures. How does learning influence chaotic systems? The macroeconomic and corporate experiments described here, for example, show the behavior of systems simulated for thousands of years and weeks, respectively. The parameters of the decision functions remain constant throughout. Yet there is evidence (Sterman 1987b, Bakken 1987) that subjects begin to learn within just a few cycles, modifying the parameters of their ordering function. In the macroeconomic experiment, at least, learning moves the subjects away from the chaotic region into the region of stability. The practical significance of chaos and other nonlinear phenomena in policy-oriented models of social and economic behavior remains clouded while these questions are unanswered.

More important, the results demonstrate the feasibility of subjecting theories of

human behavior to experimental test. Many models of human systems may be built which produce an astounding variety of behaviors. Unfortunately many of these models are formulated without regard to the vast body of empirical knowledge documenting the decision processes of individuals. For example, rational expectations models in economics presume agents have perfect knowledge of and the cognitive capability to solve the system of equations which characterize the economy (Lucas 1976, Shaw 1984). These assumptions are strongly contradicted by evidence at all levels of aggregation (Simon 1984, Hogarth and Reder 1986) and form a poor foundation for descriptive models of economic dynamics.

Similarly, physical scientists frequently specify decision processes in models of human behavior by analogy with natural systems (Beylich 1986, Zeeman 1977, Zimmerman 1986), or simply assert that the processes studied in natural systems extend by analogy to social dynamics (Ruelle 1984, May 1976). The natural sciences can indeed be a fruitful source of metaphor for theories of human behavior, and there is encouraging progress in the application of nonlinear science to social systems (e.g. Prigogine and Sanglier, 1987). But without a firm empirical grounding in the psychology of choice behavioral scientists will continue to view such models as ad hoc and continue to question their relevance. It has been difficult to assess the validity of such models because the assumed decision rules were not subject to empirical test. In the future experimental tests of models of human behavior should become as accepted and commonplace as they are in the natural sciences. Over time the application of these methods should identify the processes of judgment people actually use in dynamic decisionmaking tasks and which are acceptable as the basis for accurate modeling of human affairs. Models formulated on the basis of knowledge of individual decisionmaking, analyzed with the tools of simulation and modern nonlinear science, and subjected to experimental test offer the best hope to improve our understanding of the dynamics and evolution of social systems.

Notes

1. 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 of unfilled orders rather than as negative orders.
2. Note that the function $OK=f(\cdot)$ does not contain an estimated regression constant. Thus the correspondence of the estimated and actual capital orders, not just their variation around mean values, provides an important measure of the model's explanatory power. Since the residuals e need not satisfy $\sum e_t = 0$ the conventional R^2 is not appropriate. The alternative $R^2 = 1 - \sum e_t^2 / \sum OK_t^2$ is used (Judge et al. 1980). This R^2 can be interpreted as the fraction of the variation in capital orders around zero explained by the model.
3. 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.
4. In the experiment and in the simulations orders are restricted to the positive integers.
5. The parameters θ , α_s , β , and S' were estimated to the nearest .1, .05, .05, and 1 units, respectively. The search was carried out over a sufficiently large range to ensure capturing the global minimum of $\sum e_t^2$.

- D-5921 24
6. Note that S' functions much like a regression constant in eq. 32. However the nonlinearity $O_t \geq 0$ means the residuals will not, in general, satisfy $\sum e_t = 0$ (estimated and actual orders need not have a common mean). The conventional R^2 is not therefore 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).
 7. Subjects were assigned positions randomly, and one-way ANOVA of the estimated parameters showed no strong relationship between the values of the estimated parameters and the position one plays in the game (with the exception of α_s which was significant at the 4.3% level).

N.B. The long wave game runs on any IBM PC or compatible computer; disks are available from the author. The estimation and simulations for both experiments were carried out on Macintosh computers using the TrueBASIC language. The data and computer programs are available from the author upon request. To avoid transients, the phase plots shown in figure 4 were generated by simulating the model for 10,000 periods and discarding the first 9000. Similarly, the phase plots shown in figure 8 were simulated for 1500 weeks. The first 500 were discarded for all except the Bassbeer factory, for which the first 1000 were discarded.

References

- Bakken, B. 1987. Learning System Structure by Exploring Computer Games. unpublished manuscript, System Dynamics Group, Sloan School of Management, MIT, Cambridge MA 02139.
- Beylich, A. 1986. On the Modelling of Vehicular Traffic Flow. In C. Kilmister (ed.) *Disequilibrium and Self-Organization*. Dordrecht: D. Reidel.
- Cyert, R. and J. March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs N.J.: Prentice Hall.
- Davis, H. L., S. J. Hoch, and E. K. Easton Ragsdale. 1986. An Anchoring and Adjustment Model of Spousal Predictions. *Journal of Consumer Research*. 13:25-37.
- Day, R. 1984. Disequilibrium Economic Dynamics: A Post-Schumpeterian Contribution. *Journal of Economic Behavior and Organization*. 5, 57-76.
- Day, R. 1982a. Irregular Growth Cycles. *American Economic Review*. 72, 406-414.
- Day, R. 1982b. Complex Behavior in System Dynamics Models. *Dynamica*. 8, 82-89.
- Einhorn, H. J., and R. M. Hogarth. 1985. Ambiguity and Uncertainty in Probabalistic Inference. *Psychological Review*. 92:433-461.
- Forrester, J. W. 1961. *Industrial Dynamics*. Cambridge, MA: MIT Press.
- Forrester, J. W. and P. M. Senge. 1980. Tests for Building Confidence in System Dynamics Models. *TIMS Studies in the Management Sciences*. 14:201-228.
- Frankel, J. A. and K. A. Froot. 1987. Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations *American Economic Review* 77:133-153.
- Hines, J. 1986 A Behavioral Theory of Interest Rate Formation. Working paper WP-1771-86, Sloan School of Management, MIT.
- Hogarth, R. M. 1987. *Judgement and Choice*. 2nd ed. New York: Wiley.

- Hogarth, R. M. and M. W. Reder (eds). 1986. The Behavioral Foundations of Economic Theory *Journal of Business*. 59:S181-S505.
- Jarman, W. E. (ed.) 1963. *Problems in Industrial Dynamics*. Cambridge, MA: MIT Press.
- Johnson, E. J. and D. A. Schkade. 1987. Heuristics and Bias in Utility Assessment. Unpublished manuscript, Wharton School, University of Pennsylvania, Philadelphia.
- Judge et al. 1980. *The Theory and Practice of Econometrics*. New York: Wiley.
- Lopes, L. L. 1981. Averaging Rules and Adjustment Processes: The Role of Averaging in Inference. Report 13, Wisconsin Human Information Processing Program, University of Wisconsin, Madison.
- Lucas, R.E. 1976. Econometric Policy Evaluation: A Critique. In K. Brunner and A. Meltzer (eds). *The Phillips Curve and Labor Markets*. Supplement to the Journal of Monetary Economics.
- May, R. M. 1976. Simple Mathematical Models with Very Complicated Dynamics. *Nature*. 261(5560):459-467.
- Morecroft, J. 1983. System Dynamics: Portraying Bounded Rationality. *Omega*. 11:131-142.
- Morecroft, J. 1985. Rationality in the Analysis of Behavioral Simulation Models. *Management Science* 31:900-916.
- Mosekilde, E., D. Rasmussen, H. Jensen, J. Sturis, and J. Jespersen. 1987. Chaotic Behavior in a Generic Management Model. *European Journal of Operations Research*. Forthcoming.
- Mosekilde, E. and E. Larsen. 1987. Chaotic Behavior in the Beer Distribution Game. *System Dynamics Review*. 4(1).
- Nicolis, G. and I. Prigogine. 1977. *Self-Organization in Nonequilibrium Systems*. New York: Wiley.
- Plott, C. 1986. Laboratory Experiments in Economics: The Implications of Posted Price

- Institutions. *Science*. 232 (9 May), 732-738.
- Prigogine, I. and M. Sanglier (eds.) 1987. *Laws of Nature and Human Conduct*. Brussels: Task Force of Research Information and Study on Science.
- Rasmussen, S., E. Mosekilde, and J. D. Sterman. 1985. Bifurcations and Chaotic Behavior in a Simple Model of the Economic Long Wave. *System Dynamics Review*. 1(1):92-110.
- Richardson, G.P. and A.L. Pugh. 1981. *Introduction to System Dynamic Modeling with DYNAMO*. Cambridge MA: The MIT Press.
- Ruelle, D. 1984. Strange Attractors. In P. Cvitanovic (ed.) *Universality in Chaos*. Bristol: Adam Hilger, Ltd.
- Shaw, G.K. 1984. *Rational Expectations*. New York: St. Martins Press.
- Simon, H. A. 1979. Rational Decisionmaking in Business Organizations. *American Economic Review*. 69:493-513.
- Simon., H. A. 1984. The Behavioral and Rational Foundations of Economic Dynamics. *Journal of Economic Behavior and Organization*. 5:35-55.
- Smith, V. 1982. Microeconomic Systems as an Experimental Science. *American Economic Review*. 72, 923-955.
- Sterman, J. D. 1984. Instructions for Running the Beer Distribution Game. System Dynamics Group working paper D-3679, Sloan School of Management, MIT.
- Sterman, J. D. 1985a. A Behavioral Model of the Economic Long Wave. *Journal of Economic Behavior and Organization*. 6:17-53.
- Sterman, J. D. 1985b. A Skeptic's Guide to Computer Models. Working Paper D-3665, System Dynamics Group, Sloan Scholl of Management, MIT, Cambridge MA 02139.
- Sterman, J. D. 1986. The Economic Long Wave: Theory and Evidence. *System Dynamics Review*. 2(2):87-125.

- Sterman, J. D. 1987a. Testing Behavioral Simulation Models by Direct Experiment. *Management Science*. 33(12).
- Sterman, J. D. 1987b. Misperceptions of Feedback in Dynamic Decisionmaking. Working paper WP-1899-87. Sloan School of Management, MIT, Cambridge MA.
- Sterman, J. D. 1987c. Managerial Behavior in Dynamic Decisionmaking: Further Misperceptions of Feedback. System Dynamics Group working paper D-3919, Sloan School of Management, MIT.
- Sterman, J. D. 1987d. Expectation Formation in Behavioral Simulation Models. *Behavioral Science*. 32:190-211.
- Sterman, J. D. and D. L. Meadows. 1985. STRATEGEM-2: A Microcomputer Simulation Game of the Kondratiev Cycle. *Simulation and Games*. 16:174-202.
- Thompson, J. and H. Stewart. 1986. *Nonlinear Dynamics and Chaos*. New York: Wiley.
- Tversky, A. and D. Kahneman 1974. Judgment Under Uncertainty: Heuristics and Biases. *Science* . 185 (27 September):1124-1131.
- Zeeman, E. C. 1977. *Catastrophe Theory: Selected Papers, 1972-1977*. Reading, MA: Addison Wesley Publishing Co.
- Zimmerman, R. 1986. The Transition from Town to City: Metropolitan Behavior in the 19th Century. In C. Kilmister (ed.). *Disequilibrium and Self-Organization*. Dordrecht: D. Reidel.

Figure 1. The generic stock-management system.

Figure 2. The stock management system applied to aggregate capital investment. The model portrays a firm which produces capital plant and equipment. The demand for the firm's product is exogenous, as is the delivery delay for acquiring new capital stock. When the self-ordering multiplier is introduced, the firm's own orders for capital add to the exogenous demand, and the delivery delay for capital becomes the firm's own delivery delay.

Figure 3. Typical results of the macroeconomic experiment.

Table 1. Macroeconomic experiment: Summary of results.

	<u>Experiment</u>		<u>Optimal</u>
	Mean (Std. Dev. in parentheses)		
	N=49	N=48	
Score (units)	591 (1176)	432 (382)	19
Periodicity (years)	46 (13)	45 (11)	NO CYCLE
1st Capacity Peak (units)	2232 (1703)	3935 (1346)	630
2nd Capacity Peak (units)	1139 (671)	1139 (671)	NO 2nd PEAK
Peak Order Rate (units/period)	629 (927)	518 (501)	260
Minimum Order Rate (units/period)	4 (11)	4 (11)	0
Minimum Fraction of Demand Satisfied (%)	48 (14)	49 (13)	62

N=48 column excludes trial 1 as an outlier (score 8229; capacity peak >27,000; maximum order rate of 6000 units).

"Units" refer to units of capital

Table 2. Macroeconomic experiment: Estimated parameters.

Trial	Score	α_S	Std. Error*		α_{SL}	Std. Error*		R ²	Mode
1	8229	.88	(.04)	a	.24	(.03)	a	.90	C
2	1980	.56	(.10)	a	.26	(.13)	c	.39	P1
3	1335	1.00	(.004)	a	0.00	(.003)		.99+	C
4	1283	2.21	(.06)	a	1.51	(.06)	a	.86	C
5	965	1.35	(.01)	a	.59	(.01)	a	.99	P1
6	908	.38	(.09)	a	0.00	(.10)		.48	P1
7	890	1.28	(.02)	a	.72	(.02)	a	.97	C
8	786	.21	(.02)	a	.16	(.05)	a	.74	S
9	741	1.66	(.05)	a	1.83	(.07)	a	.76	C
10	715	.74	(.04)	a	.60	(.05)	a	.81	C
11	651	.69	(.02)	a	.33	(.03)	a	.94	C
12	638	.46	(.08)	a	.33	(.11)	a	.33	P5
13	634	3.73	(.08)	a	4.44	(.10)	a	.63	P1
14	554	.45	(.04)	a	0.00	(.06)		.87	P1
15	482	.32	(.05)	a	.20	(.10)	b	.57	S
16	479	2.49	(.04)	a	2.90	(.05)	a	.91	C
17	467	.87	(.05)	a	1.38	(.09)	a	.57	S
18	459	.62	(.04)	a	.43	(.05)	a	.86	P1
19	459	.13	(.02)	a	0.00	(.08)		.77	S
20	400	.14	(.04)	a	0.00	(.14)		.37	S
21	399	.23	(.02)	a	.28	(.06)	a	.83	S
22	384	.55	(.04)	a	0.00	(.05)		.90	C
23	373	.17	(.03)	a	.40	(.13)	a	.50	S
24	366	.56	(.06)	a	.11	(.08)		.76	C
25	311	.16	(.02)	a	0.00	(.08)		.80	S
26	267	.16	(.02)	a	0.00	(.09)		.78	S
27	265	.14	(.02)	a	0.00	(.08)		.88	S
28	241	.19	(.03)	a	0.00	(.10)		.73	S
29	229	.33	(.04)	a	0.00	(.09)		.82	P1
30	212	.22	(.02)	a	.15	(.08)	c	.85	S
31	212	.04	(.03)		.29	(.17)		.74	S
32	207	.16	(.02)	a	0.00	(.10)		.80	S
33	196	.09	(.02)	a	0.00	(.14)		.79	S
34	194	.29	(.06)	a	.25	(.13)	c	.45	S
35	189	.09	(.02)	a	0.00	(.11)		.90	S
36	183	.42	(.04)	a	0.00	(.06)		.92	P2
37	164	.14	(.04)	a	0.00	(.13)		.74	S
38	164	.22	(.02)	a	0.00	(.06)		.94	S
39	161	.02	(.02)		0.00	(.19)		.87	S
40	161	.19	(.02)	a	.02	(.07)		.93	S
41	156	.24	(.04)	a	0.00	(.11)		.82	S
42	135	.32	(.04)	a	.69	(.14)	a	.74	S
43	120	.16	(.02)	a	0.00	(.09)		.93	S
44	118	.15	(.03)	a	0.00	(.12)		.90	S
45	108	.44	(.03)	a	.40	(.07)	a	.94	S
46	105	.08	(.02)	a	0.00	(.13)		.96	S
47	100	.16	(.03)	a	0.00	(.11)		.91	S
48	90	.60	(.02)	a	.85	(.05)	a	.96	S
49	77	.11	(.03)	a	.19	(.17)		.90	S
Mean:	591	.55			.40		Pooled R ² :	.85	

*t statistic for test of $H_0: \alpha = 0$ significant at a: .01; b: .05; c: .10 level. Mode: S=Stable, Pn=period n, C=Chaotic

Figure 4. Simulation of the decision rule with estimated parameters.

Figure 5. Modes of the system in parameter space.
Lower graph zooms in to area between (0,1).

Figure 6. "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.* 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 7. Typical experimental results of the Beer distribution Game.
 7a: Customer Orders. In week 5 customer orders rise from 4 to 8 cases per week.
 Compare against the oscillations in the subjects' orders (7c).

Figure 7b. Key to experimental results (figure 7c).

Figure 7c. Typical experimental results.

Table 3. Summary of experimental results, Beer Distribution Game.
Data are averages of 11 trials.

	Total	Retailer	Wholesaler	Distributor	Factory
<hr/>					
COSTS					
Mean	\$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
Ho: Mean cost = Benchmark	p<.000+	p<.001	p<.000+	p<.000+	p<.000+
<hr/>					
	Customer	Retailer	Wholesaler	Distributor	Factory
<hr/>					
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

*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 4. Estimated parameters, Beer Distribution Game

Trial & Position		θ	α_s	β	S'	R^2	Mode
Bassbeer	R	0.90	0.10	0.65 a	20 a	0.20	S
	W	0.00	0.25 a	0.50 a	27 a	0.86	S
	D	0.15	0.05 a	0.35	14	0.74	S
	F	1.00 a	0.65 a	0.40 a	15 a	0.84	C
Budweiser	R	0.00	0.40 a	0.10 a	7 a	0.67	C
	W	0.00	0.40 a	0.75 a	30 a	0.92	S
	D	0.00	0.30 a	0.10 a	10 a	0.88	C
	F	0.25 c	0.25 a	0.10	9 a	0.87	S
Coors	R	0.00	0.20 a	0.00	25 a	0.57	S
	W	0.00	0.15 a	0.50 a	38 a	0.11	S
	D	0.90 a	0.30 a	0.20 a	10 a	0.61	S
	F	0.25	0.30 a	0.00	18 a	0.73	P12
Freebeer	R	0.40	0.35 a	0.45 a	15 a	0.43	S
	W	0.30	0.05 a	0.00	30 c	0.76	S
	D	0.05	0.35 a	1.00 a	18 a	0.86	S
	F	0.25	0.25 a	0.00	19 a	0.89	P1
Grin & Beer It	R	0.10	0.35 a	0.65 a	13 a	0.60	S
	W	0.95 a	0.15 a	0.55 a	14 a	0.79	S
	D	0.20 b	0.20 a	0.30 a	19 a	0.94	S
	F	0.25	0.35 a	0.55 a	24 a	0.73	S
Grizzly	R	0.05	0.30 a	0.65 a	31 a	0.58	S
	W	0.30	0.20 a	0.35 a	27 a	0.82	S
	D	0.15	0.05	0.25	15	0.32	S
	F	0.55 a	0.65 a	0.00	9 a	0.75	C
Heineken1	R	0.95	0.15 a	0.00	9 a	0.75	S
	W	0.50 a	0.00	N/D	N/D	0.87	S
	D	0.20 a	0.30 a	0.05 a	8 a	0.98	S
	F	0.80 b	0.00	N/D	N/D	0.60	S
Heineken2	R	0.50	0.05	0.60	6	0.10	S
	W	0.40 a	0.10 a	0.30 a	16 a	0.81	S
	D	1.00 a	0.15 a	0.80 a	14 a	0.73	S
	F	0.55 a	0.80 a	0.00	9 a	0.87	C
Heineken3	R	0.05	0.30 a	0.45 a	5 a	0.89	P5
	W	0.20	0.00	N/D	N/D	0.23	S
	D	0.30 a	0.10 a	0.90 a	12 a	0.94	S
	F	0.00	0.30 a	0.15 c	17 a	0.87	C
Suds	R	1.00	0.00	N/D	N/D	0.76	S
	W	0.05	0.30 a	0.20 a	20 a	0.76	S
	D	0.15	0.60 a	0.35 a	0	0.69	C
	F	0.40 a	0.35 a	1.05 a	32 a	0.95	S
Twoborg	R	0.75	0.35 a	0.00	4 a	0.83	C
	W	0.00	0.25 a	0.05	18 a	0.72	C
	D	0.05	0.50 a	0.00	15 a	0.84	C
	F	0.95 a	0.30 b	0.20	26 a	0.66	S
Minimum		0.00	0.00	0.00	0	0.10	
Maximum		1.00	0.80	1.05	38	0.98	
Mean		0.36	0.26	0.34	17	0.71	

N/D: Not Defined

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

Mode: S=Stable, Pn=Limit cycle of period n, C=Chaotic

Figure 8. Simulation of decision rule with estimated parameters, Beer Distribution Game.

Figure 9. Beer Distribution Game. Modes of the system in parameter space.

1813 0810

Date Due

JY 25

SEP 15 1988

NOV 15 1990

FEB 27 1992

NOV. 08 1995

MIT LIBRARIES

3 9080 004 928 328