Experimental Evidence of Deterministic Chaos in Human Decision Making Behavior

John D. Sterman*, Erik Mosekilde†, Jesper Skovhus Thomsen†

*MIT Sloan School of Management, Cambridge, MA 02139

†Physics Laboratory III, Technical University of Denmark, DK-2800 Lyngby, Denmark

Sloan Working Paper No. BPS-3241-91 January 1991

D-3957-2

Experimental Evidence of Deterministic Chaos in Human Decision Making Behavior

John D. Sterman*, Erik Mosekilde†, Jesper Skovhus Thomsen†

January 1991

* Massachusetts Institute of Technology, Sloan School of Management, Cambridge, MA 02139

† Physics Laboratory III, Technical University of Denmark, DK-2800 Lyngby, Denmark

Summary

An experiment with a simulated economy demonstrates that the decision-making processes of human subjects can produce deterministic chaos. Participants managed a commodity production-distribution system to minimize costs. Performance, however, was systematically suboptimal. Econometric estimation of subjects' decision rules identifies the sources of poor performance. Simulation of the estimated rules reveals nonlinear phenomena including chaos, hyper-chaos, quasiperiodicity, mode-locking, and coexisting stable and unstable solutions. The results show the applicability and importance of modern nonlinear dynamics in models of human systems.

The prevalence of deterministic chaos and other nonlinear phenomena in physical, chemical, and biological systems has prompted speculation that these dynamics may occur in human systems as well[1-3]. Indeed, numerous models have shown how economic systems might produce chaos[4-7]. But the significance of these theoretical developments hinges on whether the chaotic regimes lie in the realistic region of parameter space. Further, the decision rules postulated in these models have not been tested experimentally. Despite intriguing efforts to identify chaos in economic time series[8,9], it is difficult to resolve such issues by empirical means alone[10,11]. Data series are often unavailable, or too short relative to the frequencies of interest. Measurement error and process noise in social and economic data further complicate empirical analysis. A complementary approach is based on laboratory experiments in which models provide a simulated environment for the study of decision-making[12]. We report here the results of one such experiment which demonstrates that the decision-making processes of human subjects can produce deterministic chaos in a common economic context.

The experiment simulates an industrial production-distribution system. Such systems offer firms flexibility through a decentralized network of inventories which buffer differences between the demand for and production of goods. Production-distribution systems have long been associated with business cycles and other fluctuations in industrial economies[13,14]. The experiment here, the 'Beer Distribution Game,' is a role-playing simulation of a typical production-distribution system. Developed at MIT to introduce students of management to economic dynamics and computer simulation, the game has been widely used for thirty years[15].

The experiment is conducted on a board which portrays in a simplified fashion the production and distribution of beer (Fig. 1). Beer is represented by markers which are manipulated by the players as the game proceeds. Each brewery consists of four sectors: retailer, wholesaler, distributor, and factory (R, W, D, F). One subject manages each sector. Customer demand is represented by a deck of cards. Each simulated week customers order beer from the retailer, who ships the beer requested out of inventory. The retailer in turn orders from the wholesaler, who likewise ships out of inventory. Similarly, the wholesaler orders and receives beer from the distributor, who orders and receives beer from the factory. The factory brews the beer. At each stage there are shipping and order receiving delays. Each week the subjects must decide how much to order from their immediate supplier (factories decide how much beer to brew).

Subjects seek to minimize cumulative costs over 36 simulated weeks. Inventory holding costs for each sector are $.50/case/week, and backlogs (negative inventories) cost $1.00/case/week. Subjects must therefore keep their inventory low while avoiding backlogs. Due to the order receiving and shipping lags, a substantial supply line of unfilled orders may build up. The lag in receiving beer may vary: if the wholesaler can cover the retailer's orders, the retailer quickly receives the beer desired. But if the wholesaler has run out, the retailer must wait until the wholesaler can restock. Although the experimental system is simplified, it nevertheless captures many features of real economies, including multiple feedbacks, nonlinearities, and time lags. Nonlinearities arise from nonnegativity constraints on orders and shipments: shipments normally equal incoming orders, but if inventory is depleted, shipments must equal zero, and the unfilled orders remain in the backlog for future delivery. The production-distribution structure has been validated for a variety of industries[13,16,17].

The experiment follows standard protocols[18,19] and is detailed in ref. 20. Subjects were graduate and undergraduate students at MIT and senior executives from a number of U.S. firms (N=44). Each trial begins in equilibrium. Each inventory contains 12 cases. Customer demand is initially 4 cases/week. Equilibrium is disturbed by an unannounced step increase in customer demand to 8 cases/week in week 5.

Results are summarized in Table 1; Fig. 2 shows a typical trial. The participants' performance is significantly suboptimal. Total costs averaged $2028, ten times higher than the optimal costs of $204, a highly significant difference ($t = 8.7$, $p<.00001$). More interesting, the results exhibit strong regularities, suggesting subjects used a common heuristic in ordering: *1. Oscillation:* Orders and inventories exhibit a large amplitude fluctuation. On average 21 weeks are required to recover initial inventory levels. *2. Amplification:* The variance of orders rises steadily from customer to retailer to factory. Customer orders increase from 4 to 8 cases/week; responding to this disturbance, factory orders rise to a peak averaging 32 cases. *3. Phase lag:* Orders tend to peak later as one moves from retailer to factory. All three characteristics are well documented in industrial economies[13,14,16,21].

We next estimate a model of the subjects' decision rule[20]. Demand uncertainty and large backlog costs mean subjects should maintain a small but positive inventory. To do so subjects should order enough to (i) meet expected demand D^e, (ii) correct discrepancies between the desired stock S' and actual stock S, and (iii) ensure a steady flow of deliveries by maintaining an adequate supply line of unfilled orders SL. (The supply line is the accumulation of orders placed but not yet received.) Since order cancellations are not allowed, orders O must be nonnegative, yielding

$$O_t = \mathrm{MAX}(0, D^e_t + \alpha(S' - S_t - \beta SL_t)) \tag{1}$$

where α is the fraction of stock discrepancies corrected each period and β is the fraction of the supply line subjects consider. Adaptive expectations are assumed for each subject's forecast of incoming orders (exponential smoothing of demand D):

$$D^e_t = \theta D_{t-1} + (1-\theta)D^e_{t-1}, \quad 0 \le \theta \le 1. \tag{2}$$

The adjustment term $\alpha(S' - S_t - \beta SL_t)$ creates two negative feedback loops which regulate inventories. Discrepancies between the desired and actual stock induce additional orders until inventory reaches the desired value. Orders also slow once sufficient orders to restore inventory have been placed in the supply line – depending on β. If β=0, then orders placed are ignored until they arrive, causing overordering and instability. If β=1, then subjects fully account for the supply line and do not double order. While β=1 is optimal, a prior experiment showed β<<1 for many subjects in a similar task[22]. Consistent with behavioral decision theory[23-25], the rule utilizes information locally available to the decision maker and does not presume managers have the cognitive capability to solve for optimum performance.

The parameters were estimated for each participant by nonlinear least squares subject to the constraints $0 \leq \theta \leq 1$ and α, β, $S' \geq 0$ (table 2). The explanatory power of the proposed rule is excellent: the mean R^2 is 71%; R^2 is less than 50% for only 6 of 44 subjects. A large majority of the estimated parameters are significant, and the estimated parameters are systematically related to performance[20].

Though the subjects' disequilibrium response is suboptimal, we expected that their decision rules would be stable and swiftly return the system to a low-cost equilibrium. To test this hypothesis we simulated the experimental system using each set of estimated parameters. Analysis of variance showed no strong relations between subjects' position in the distribution chain and the parameters of their decision rule. We therefore assume identical parameters for each sector, reducing the dimension of the parameter space from 16 to 4. Thirty parameter sets (68%) do indeed produce stable behavior. However, one periodic and three quasiperiodic solutions appear. Ten (23%) yield chaotic behavior. The sizes and shapes of the chaotic attractors produced by simulation of the parameters characterizing different subjects vary widely (Fig. 3): modest changes in cue weights produce large changes in dynamics. Though the experiment is a difference equation system, the continuous-time analog yields similar chaotic dynamics[26].

To explore the structure of the parameter space, we set θ and S' at representative values of .25 and 17, respectively, and varied α and β over the interval [0,1]. The space (Fig. 4) includes stable, periodic, quasiperiodic, and chaotic solutions. The unstable solutions arise from the nonlinear coupling of several oscillatory feedbacks created by the multiple inventories and time delays in the system. These oscillators are coupled nonlinearly through the availability of inventory in the distribution chain. High frequencies are produced when there is sufficient inventory so that orders can be filled by each sector's immediate supplier, e.g., when the retailer's orders are filled out of the wholesaler's inventory. Low frequencies arise when inventories are inadequate, forcing downstream sectors to wait for the factory to receive, produce, and ship the orders.

In general, larger values of α and smaller values of β are destabilizing (Fig. 4a). To the extent a subject ignores the supply line of unfilled orders ($\beta<1$), a stock shortfall causes 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 aggressive stock adjustment (larger α) since more is ordered in response to a given stock shortfall. However, the boundaries between modes are not simple. Note particularly the narrow fjords of stable solutions which snake between the periodic and chaotic solutions.

Fig. 4b magnifies a region of parameter space at the transition from stable to unstable solutions. The alternating bands of periodic and aperiodic solutions arise from mode-locking between the various sectors and resemble the Arnol'd tongues associated with a devil's staircase[27,28]. The structure of the phase diagram is further complicated, however, by the fingers of stable solutions which cut across the tongues. This complexity is caused by the coexistence of stable and unstable solutions for the same parameter values; initial conditions determine which solution is realized.

The Lyapunov spectrum ranks the system's Lyapunov exponents from largest to smallest and indicates the steady-state mode of behavior. Positive exponents indicate that, on average, nearby trajectories diverge and imply the system is chaotic; multiple positive exponents indicate divergence in multiple dimensions of phase space, a phenomenon called hyperchaos. Negative exponents indicate local convergence and imply stable or periodic behavior. Fig. 5 shows large regions of parameter space exhibit higher-order hyperchaos (three positive exponents, indicating nearby trajectories diverge along three dimensions of phase space), a behavior only rarely seen, and to our knowledge, never before in a human system. Significantly, α and β for many of the subjects fall in the regions of simple and higher-order chaos. The chaotic solutions exist in the managerially meaningful region of parameter space.

It is common in the social sciences to assume that decision-making behavior and thus the dynamics of human systems are, if not optimal, then at least stable. These results show that formal rules which characterize actual managerial decision making can produce an extraordinary range of disequilibrium dynamics, including chaos, mode-locking, coexisting stable and unstable solutions, and other highly nonlinear phenomena. Though complex relative to prior models of chaotic dynamics, the experimental system portrays in a simple but realistic manner a structure found in all modern economies. The experimental subjects, including experienced managers, produce costly fluctuations similar to those observed in reality. Such complexity raises important issues for social scientists. Policy interventions often imply changes in the parameters of a decision rule or model. But if the 'policy space' contains fractal boundaries, changes on the margin may produce unpredictable qualitative changes in behavior. Experience may not transfer to circumstances which differ only slightly. Do robust principles of policy design exist in such systems? Does chaos slow the discovery of cause and effect by agents in the economy and thus hinder learning or evolution towards efficiency? Indeed, does learning alter the parameters of decision rules so that systems evolve towards or away from the chaotic regime? While further development of theory and experiment are required to answer these questions, the results show the importance and feasibility of analyzing complex social behavior with the tools of modern nonlinear theory.

1. May, R. M. *Nature* **261**, 459-467 (1976).
2. Thompson, J. & Stewart, H. *Nonlinear Dynamics and Chaos* (Wiley, Chichester, 1986).
3. Cvitanovic, P. *Universality in Chaos* (Hilger, Bristol, 1984).
4. Day, R. *Am. Econ. Rev.* **72**, 406-414 (1982).
5. Baumol, W. & Benhabib, J. *J. Economic Perspectives* **3**, 77-105 (1989).
6. Grandmont, J. M. *Econometrica* **53**, 995-1046 (1985).
7. Lorenz, H. W. *Nonlinear Dynamical Economics and Chaotic Motion* (Springer Verlag, Berlin, 1989).
8. Chen, P. *System Dynamics Review* **4**, 81-108 (1988).
9. Candela, G. & Gardini, A. *J. of Econ. Dynamics and Control* **10**, 249-254 (1986).
10. Brock, W. A. *J. Econ. Theory* **40**, 168-195 (1986).
11. A. Wolf, J. Swift, H. Swinney, J. Vastano, *Physica* **16D**, 285-317 (1985).
12. Sterman, J. D. *Management Science* **33**, 1572-1592 (1987).
13. Forrester, J. W. *Industrial Dynamics* (MIT, Cambridge, 1961).
14. Zarnowitz, V. *Orders, Production, and Investment - Cyclical and Structural Analysis* (NBER, New York, 1973).
15. Jarmain, W. E. *Problems in Industrial Dynamics* (MIT, Cambridge, 1963).
16. Mass, N. *Economic Cycles: An Analysis of Underlying Causes* (MIT, Cambridge, 1975).
17. Senge, P. *Socioeconomic Planning Sciences* **14**, 269-280 (1980).
18. C. Plott, *Science* **232**, 732 (1986).
19. V. Smith, *Am. Econ. Rev.* **72**, 923 (1982).
20. Sterman, J. D. *Management Science* **35**, 321-339 (1989).
21. Hansen, A. *Business Cycles and National Income* (Norton, New York, 1951).
22. Sterman, J. D. *Org. Behavior and Human Decision Processes* **43**, 301-335 (1988).
23. Hogarth, R. *Judgement and Choice,* 2nd ed. (New York, Wiley, 1987).
24. Tversky, A. & Kahneman, D. *Science* **185**, 1124-1131 (1974).
25. Simon, H. A. *Am. Econ. Rev.* **69**, 493-513 (1979).

26. Mosekilde, E. & Larsen, E. *System Dynamics Review* **4**, 131-147 (1988).

27. Jensen, M., Bak, P. & Bohr, T. *Phys. Rev. A* **30**, 1960-1969 (1984).

28. Glazier, J., Jensen, M., Libchaber, A & Stavans, J. *Phys. Rev. A* **34**, 1621-1624 (1986).

TABLE 1 Summary of experimental results.

	Customer	Retailer	Wholesaler	Distributor	Factory
TIMING					
Mean time to recover initial inventory (weeks)	N/A	24	23	22	16
AMPLIFICATION					
Mean Peak Order Rate (cases/week)	8	15	19	27	32
Standard Deviation of Order Rate (cases/week)	1.6	3.6	4.8	6.7	8.5
PHASE LAG					
Mean Date of Peak Order Rate (week)	5	16	16	21	20

TABLE 2 Summary of estimation results (N=44)

Parameter:	θ	α	β	S'	R^2
Mean (std. dev.)	0.36 (0.35)	0.26 (0.18)	0.34 (0.31)	17 (9)	0.71 (0.22)
Median:	0.25	0.28	0.30	15	0.76
Minimum:	0.00	0.00	0.00	0	0.10
Maximum:	1.00	0.80	1.05	38	0.98

The ordering heuristic is given by eqs. 1-2. θ determines the speed of adjustment of the demand forecast; α determines the response to inventory shortfalls; β is the fraction of the supply line accounted for by the subjects; S' is the desired stock[20].

FIG 1. 'Beer Distribution Game' board. Each simulated week, subjects: 1. Receive inventory and advance shipping delays; 2. Fill orders; 3. Record inventory or backlog; 4. Advance the incoming orders; 5. Place orders. Subjects make their decisions in step 5; steps 1-4 handle mechanics and bookkeeping.

FIG 2. Typical experimental results. Top: customer orders increase from 4 to 8 cases/week in week 5. Middle: The resulting orders placed by subjects (from bottom to top, Retailer, Wholesaler, Distributor, Factory). Bottom: Effective inventory levels (Effective inventory = Inventory - Backlog). Tick marks on y-axes denote 10 units. Note the oscillation, amplification, and phase lag as the disturbance propagates from customer to factory.

FIG 3. Simulation of the decision rule with estimated parameters. Top: phase portrait showing retailer inventory vs. wholesaler inventory for parameters of subject 4 (θ, α, β, S' = 1.0, 0.65, 0.40, 15). Middle: the same variables for system simulated with parameters of subject 21 (θ, α, β, S' = 0.55, 0.65, 0, 9). Both are 18,000 week simulations with first 8,000 periods deleted to remove transient; flow is generally clockwise. Bottom: The parameters for subject 27 (θ, α, β, S' = 0.2, 0.3, 0.05, 8), though stable, produce a long chaotic transient (Retailer orders shown).

FIG 4. Distribution of modes in the (α,β) plane. *a,* The stock and supply line adjustment parameters α and β are varied over the interval [0,1] in increments of .005, for $\theta = 0.25$ and $S' = 17$. Note the fjords of stable solutions separating regions of periodic and aperiodic behavior. *b,* 10x magnification of the region $0.35 \leq \alpha \leq 0.42$, $0.02 \leq \beta \leq 0.12$ (outlined area in *a*). Note the complex distribution of periodic and aperiodic modes and the fingers of stable behavior which penetrate the region of unstable behavior, indicating coexisting solutions.

FIG 5. Map of Lyapunov spectrum in the region $0 \leq \alpha,\beta \leq 1$ for the same values of θ and S' as in Fig. 4, showing signs of the three largest Lyapunov exponents for each point. The region bounded approximately by $\alpha > .5$ and $\beta < .5$ contains modes with three positive Lyapunov exponents, indicating higher-order hyperchaos.

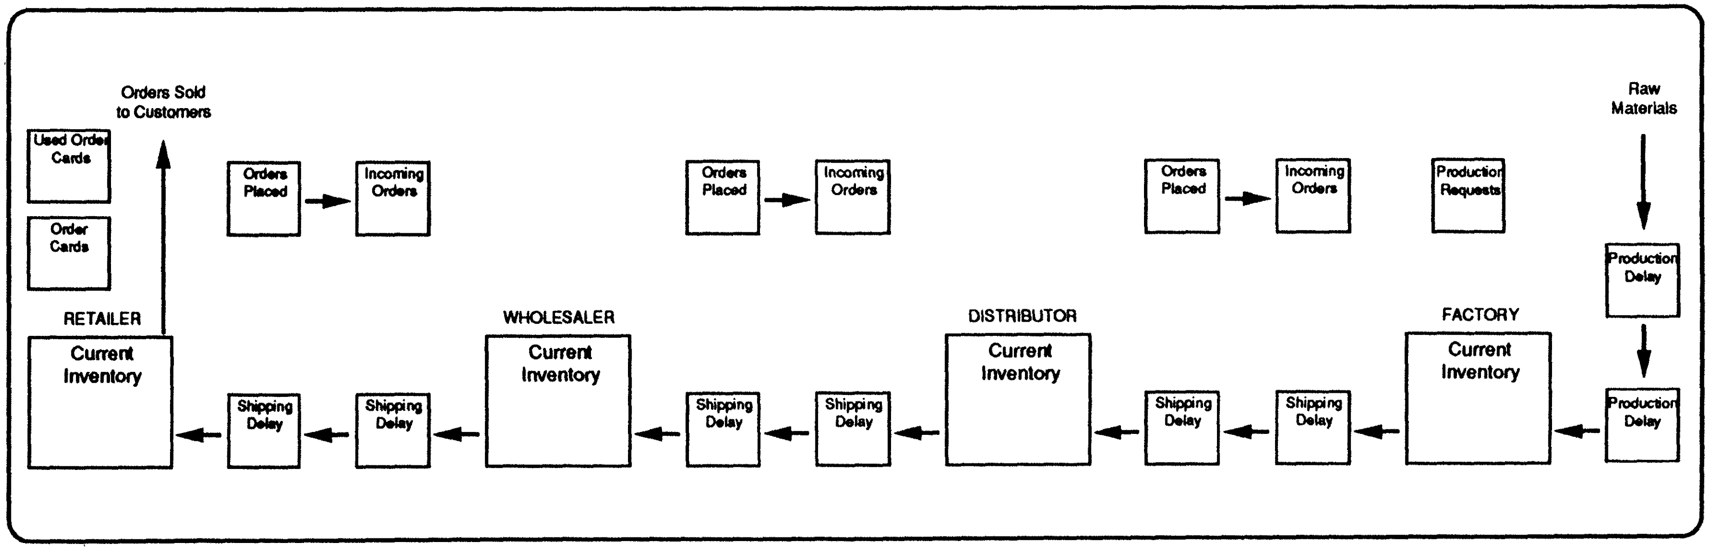
Orders Sold
to Customers
Used Order Cards
Order Cards
RETAILER
Current Inventory
Orders Placed
Incoming Orders
Shipping Delay
Shipping Delay
WHOLESALER
Current Inventory
Orders Placed
Incoming Orders
Shipping Delay
Shipping Delay
DISTRIBUTOR
Current Inventory
Orders Placed
Incoming Orders
Shipping Delay
Shipping Delay
Production Requests
FACTORY
Current Inventory
Raw Materials
Production Delay
Production Delay

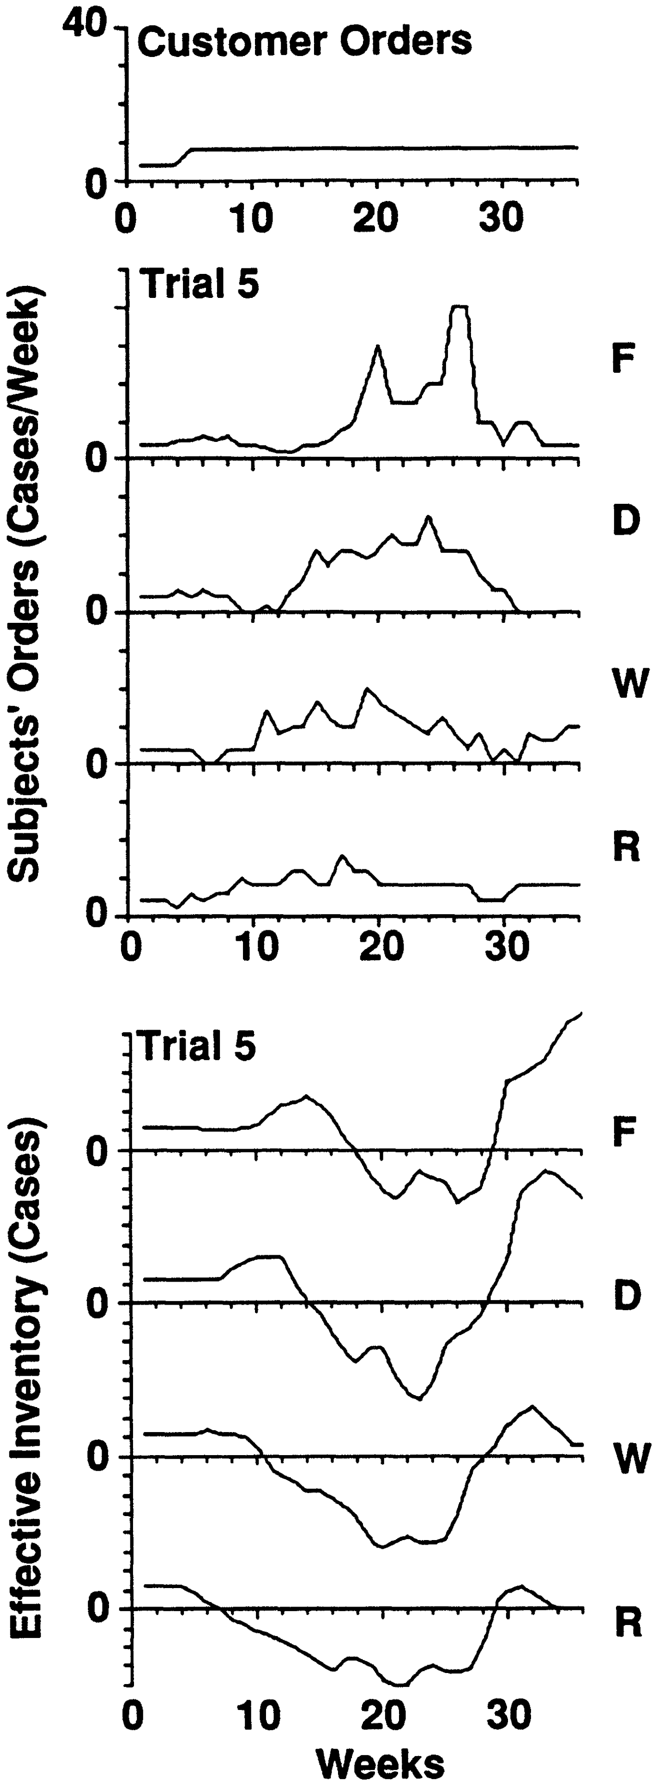

Customer Orders
40
0
0
10
20
30
Trial 5
Subjects' Orders (Cases/Week)
F
D
W
R
Trial 5
Effective Inventory (Cases)
Weeks

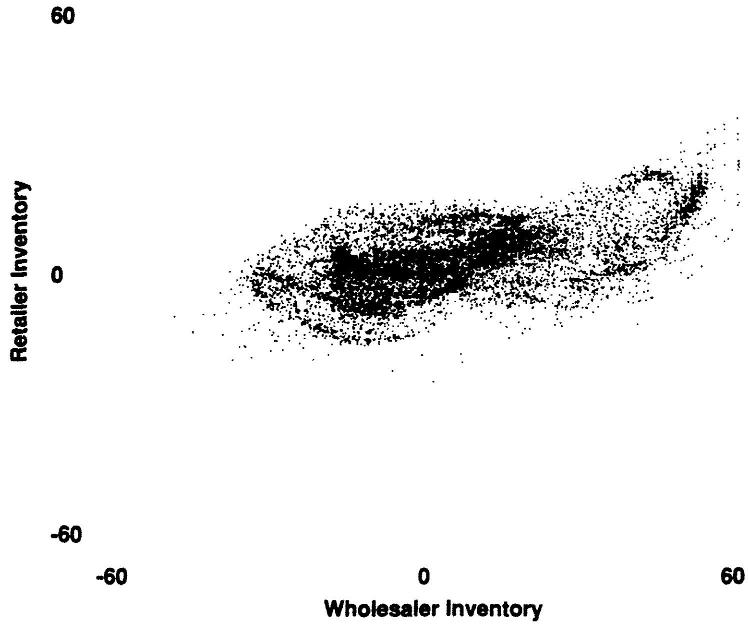
60
0
-60
Retailer Inventory
-60
0
60
Wholesaler Inventory

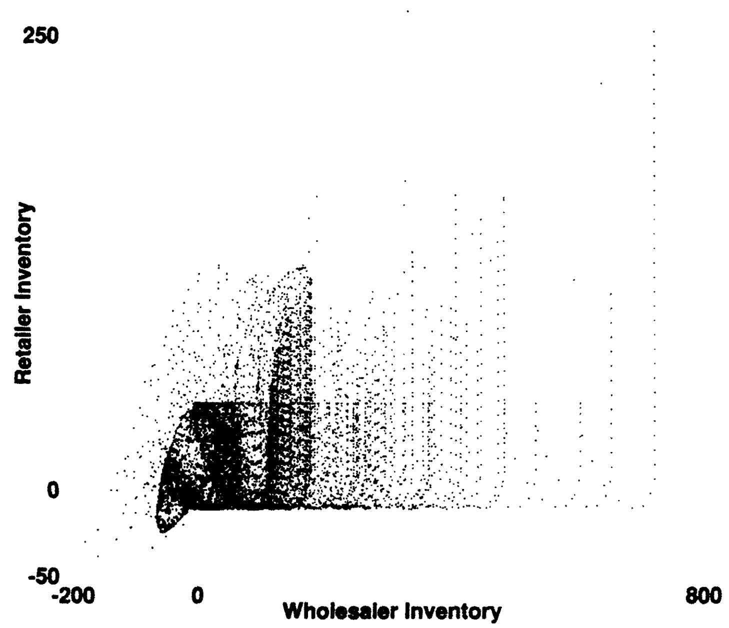
250
0
-50
Retailer Inventory
-200
0
800
Wholesaler Inventory

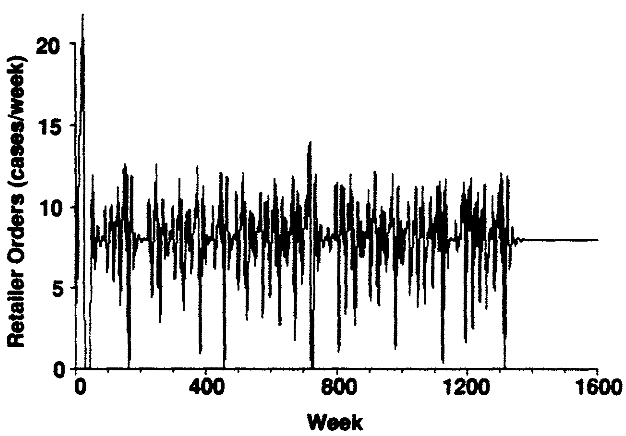
20
15
10
5
0
Retailer Orders (cases/week)
0
400
800
1200
1600
Week

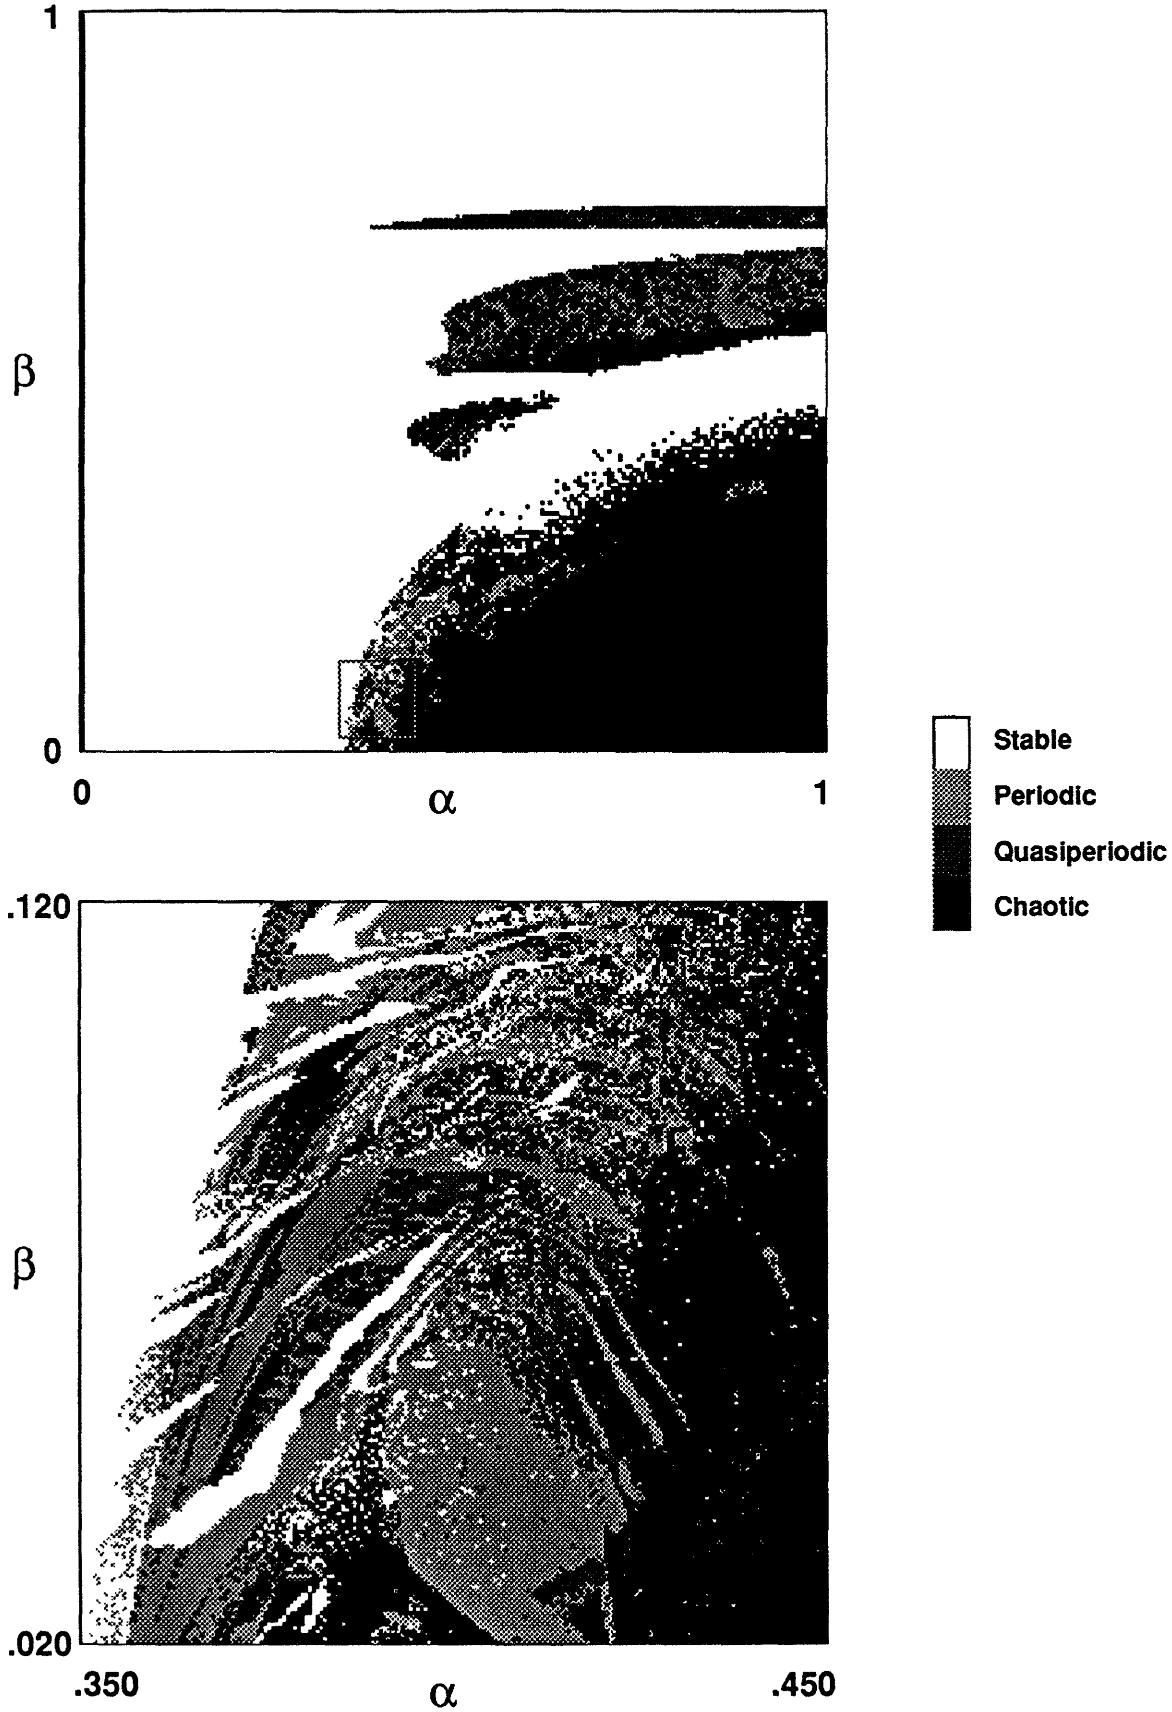
1
0
β
0
α
1
Stable
Periodic
Quasiperiodic
Chaotic
.120
.020
β
.350
α
.450

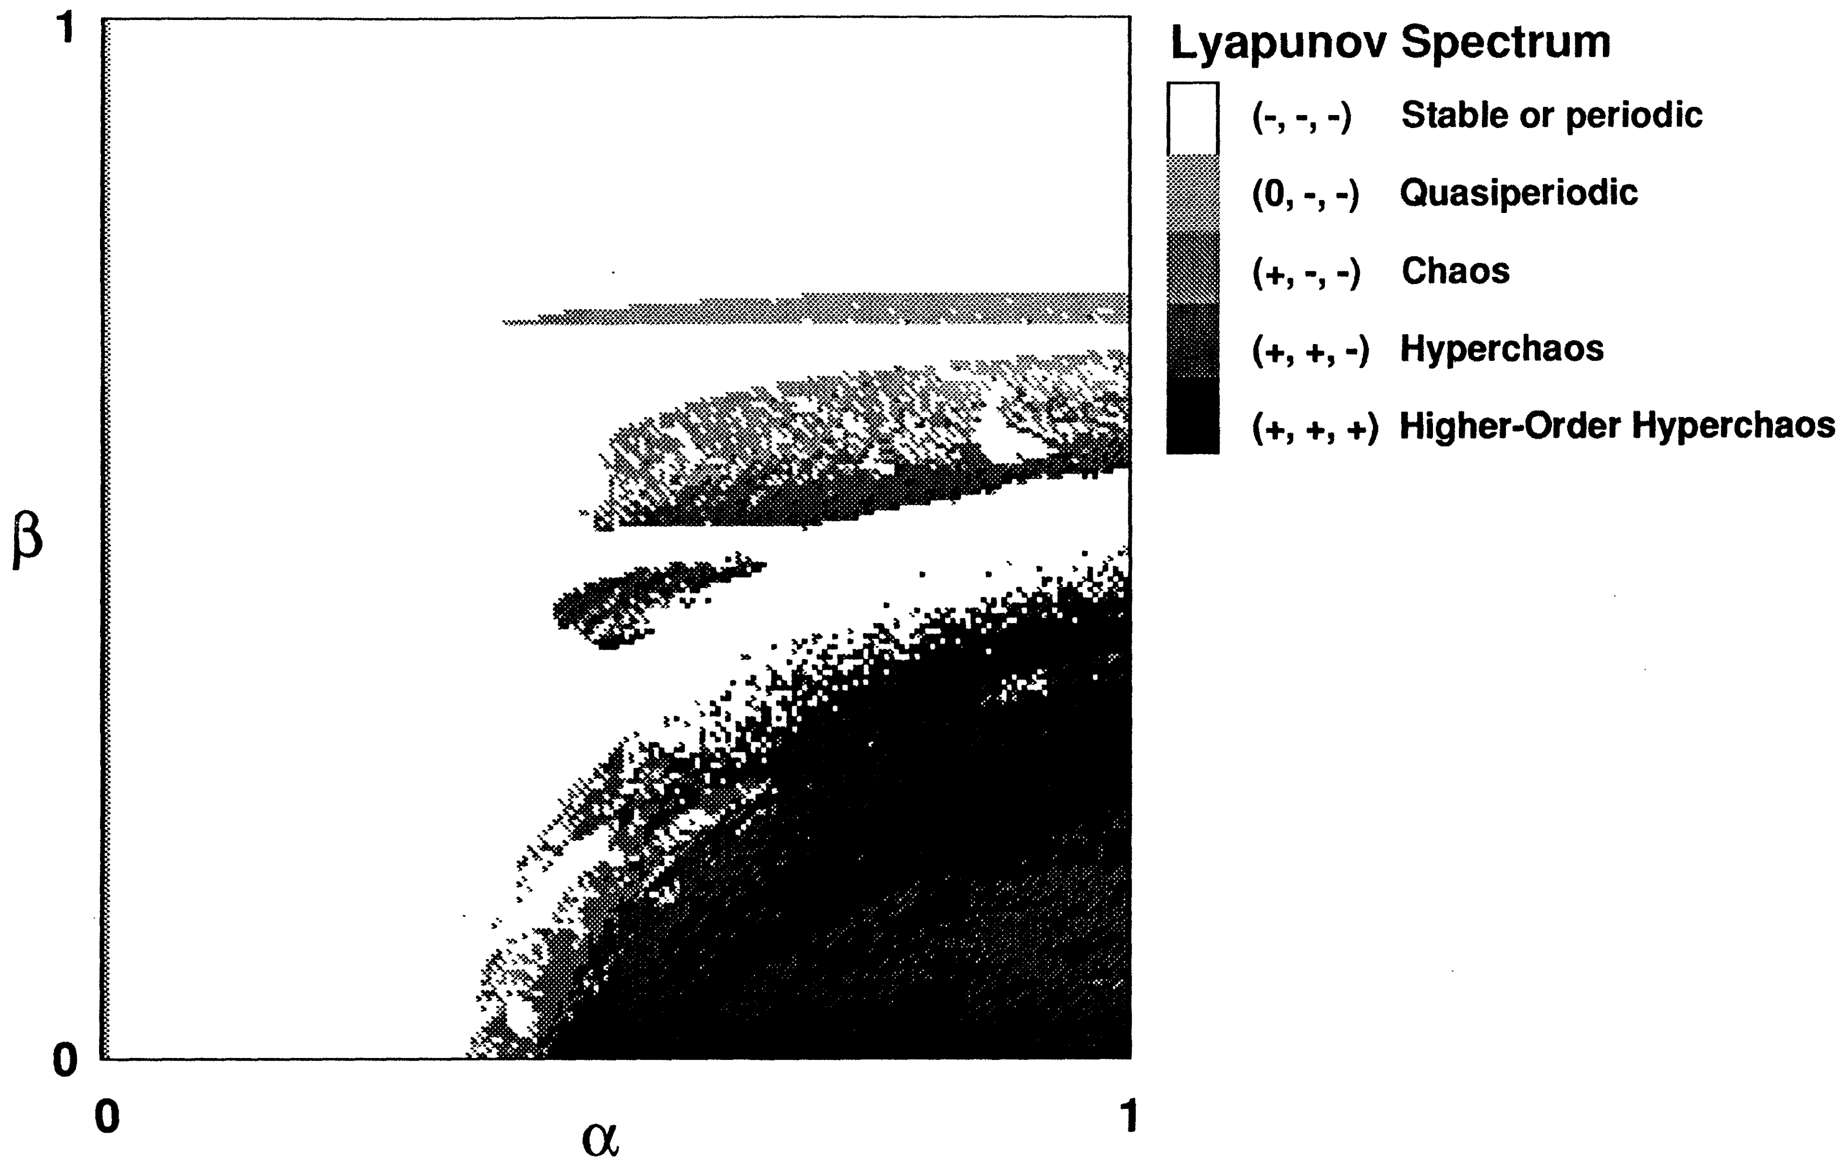
1
0
β
0
1
α
Lyapunov Spectrum
(-, -, -) Stable or periodic
(0, -, -) Quasiperiodic
(+, -, -) Chaos
(+, +, -) Hyperchaos
(+, +, +) Higher-Order Hyperchaos

Appendix: Equations for Production-Distribution System

The equations are presented in the sequence of calculation used in the experiment and simulations (Fig. 1). Initial equilibrium conditions are: all inventories = 12 units, all backlogs = 0, all other variables = 4 units (orders placed, incoming orders, production delays, shipping delays, and expected orders). Customer demand is initially 4 and rises to 8 in week 5. Subscripts denote the sector of the distribution chain (Retailer R, Wholesaler W, Distributor D, and Factory F). The computer programs for the simulations and parameter estimation are available from the first author.

Step 1: The contents of the shipping delay (D1) immediately to the right of each sector's inventory (I) are added to inventory. The contents of the shipping delay on the far right (D2) are moved into D1. Factories advance the production delays $D1_F$ and $D2_F$ in the same fashion.

$$I_j(t) = I_j(t-1) + D1_j(t-1) \tag{A1}$$

$$D1_j(t) = D2_j(t-1) \tag{A2}$$

Step 2: Retailers examine the top card on the Customer Order deck (CO); all others examine Incoming Orders (IO). All sectors fill orders. The Shipment rate S must equal the new orders plus any Backlog B from the prior period, to the extent inventory permits. The retailer's shipments go directly to the customer and leave the system. Shipments of all others are placed in the shipping delay D2 of the downstream sector (A4). Shipments also reduce inventory (A5).

$$S_R(t) = \mathrm{MIN}(CO(t) + B_R(t-1), I_R(t)) \tag{A3}$$

$$S_l(t) = \mathrm{MIN}(IO_{l-1}(t-1) + B_l(t-1), I_l(t)), \quad l = W, D, F \tag{A3'}$$

$$D2_k(t) = S_{k+1}(t), \quad k = R, W, D \tag{A4}$$

$$I_j(t) = I_j(t) - S_j(t) \tag{A5}$$

Step 3: All sectors record inventory or backlog. The net change in backlog is the difference between incoming orders and shipments. Effective Inventory EI is inventory less backlog.

$$B_R(t) = B_R(t-1) + CO(t) - S_R(t) \tag{A6}$$

$$B_l(t) = B_l(t-1) + IO_{l-1}(t-1) - S_l(t), \quad l = W, D, F \tag{A6'}$$

$$EI_j(t) = I_j(t) - B_j(t) \tag{A7}$$

Step 4: Each sector advances the slip containing last week's Order O to incoming orders. The factory puts last week's order in the top production delay $D2_F$.

$$IO_k(t) = O_k(t-1), \quad k = R, W, D \tag{A8}$$

$$D2_F(t) = O_F(t-1) \tag{A8'}$$

Step 5: Each sector places orders. First the Supply Lines SL are calculated. The supply line is the sum of units in the two shipping delays, the backlog of the supplier (if any) and orders placed the previous week. Since the factory is the primary producer, its supply line is simply the contents of the production delays.

$$SL_k(t) = D1_k(t) + D2_k(t) + B_{k+1}(t) + IO_k(t), \quad k = R, W, D \tag{A9}$$

$$SL_F(t) = D1_F(t) + D2_F(t) \tag{A9'}$$

Each sector's forecast of incoming orders (expected Demand) D^e is formed by adaptive expectations, with smoothing parameter θ.

$$D_j^e(t) = \theta_j IO_j(t-1) + (1-\theta_j) D_j^e(t-1) \tag{A10}$$

Finally, orders are the given by demand forecast adjusted by a fraction α of the difference between the desired Stock S' and the effective inventory, including a fraction β of the supply line. Orders must be nonnegative.

$$O_j(t) = MAX(0, D_j^e(t) + \alpha_j(S'_j - EI_j(t) - \beta_j SL_j(t))) \tag{A11}$$