D-3783

TESTING BEHAVIORAL SIMULATION MODELS BY DIRECT EXPERIMENT

by

John D. Sterman

WP-1752-86

D-3783

TESTING BEHAVIORAL SIMULATION MODELS BY DIRECT EXPERIMENT

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

February 1986

D-3783

TESTING BEHAVIORAL SIMULATION MODELS BY DIRECT EXPERIMENT

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

Abstract

The decision rules in simulation models purport to describe decision-making behavior as it is and not as it should optimally be. Without the criterion of optimality to judge the appropriateness of a decision rule, simulation modelers must rely on empirical confirmation of the structure of their models. In models of small organizations, traditional social science methods may be used. But these methods are infeasible in models of larger systems such as industries or the macroeconomy. This paper shows how direct experiment can be used to confirm or disconfirm the decision rules in simulation models. Direct experiment uses interactive gaming in which human subjects play a role in the system being modeled. The subjects play the game in the same physical and institutional context assumed in the model, and are given the same information set, but are free to make decisions any way they wish. The behavior of the subject can then be directly compared against the behavior produced by the assumed decision rules of the model. An example is described in detail and the correspondence of the experiment to reality is discussed.

TESTING BEHAVIORAL SIMULATION MODELS BY DIRECT EXPERIMENT

The problem of testing behavioral simulation models

The utility of simulation models depends on the confidence the model users vest in the model. The model must represent the physical and institutional structure of the system and the decisionmaking procedures used by the actors with enough accuracy for the purpose at hand. Accurately portraying the 'physics' of the system is relatively straightforward. In contrast, discovering and representing the decision rules of the actors is subtle and difficult. In models of small organizations such as a family, community, or corporation, traditional social science techniques can be used to gather primary data on decisionmaking behavior. Interviews, surveys, participant observation, and other techniques can reveal the networks of information flow, organizational structures, and decisionmaking heuristics necessary to construct a useful model.

Such techniques are of less use to the analyst interested in larger systems such as an entire industry or the macroeconomy. Fieldwork involving a significant sample of firms is prohibitively expensive and time consuming. Consistent aggregation is difficult. The traditional alternative has been to draw on established organizational and economic theory to specify the model, followed by estimation of the parameters and sensitivity tests. Econometric estimation provides an obvious means to test the consistency of models with past experience.

But these methods are unsatisfying to many economists and simulation modelers alike. Data limitations and technical difficulties of identification and estimation aside, econometrics is fundamentally unable to validate the behavioral decision rules in simulation models because the data represent the 'what' of decisions, not the 'why'. The numerical data used in estimation are the result of decisionmaking, and do not in themselves reveal the motivation for the decisions. As a result, econometrics has proven to be a rather dull knife: it is often impossible to discriminate between radically divergent theories using econometrics alone (Leamer 1983, Thurow 1983, Leontief 1982, 1971, Phelps-Brown 1972, Keynes 1939).

Of more importance, however, traditional neoclassical economic theory is heavily based on the assumptions of rational behavior, optimization, and equilibrium. Human behavior is assumed to be rational: decisions are guided by the urge to maximize profits or utility; the information required to successfully optimize is assumed to be available, usually freely, and often including information about the true structure of the system (as in rational expectations), about the future (as in intertemporal optimization models) and about hypothetical situations (e.g. the productivity of untried combinations of factor inputs). The economy is assumed to be in or near equilibrium nearly all the time, and adjustment processes are usually assumed to be stable.[1]

The behavioral simulation modeler cannot accept such assumptions. As Herbert Simon (1979, 510) declared in his Nobel Prize acceptance speech,

> There can no longer be any doubt that the micro assumptions of the theory--the assumptions of perfect rationality--are contrary to fact. It is not a question of approximation; they do not even remotely describe the processes that human beings use for making decisions in complex situations.

The purpose of simulation models is to mimic the real system so that its behavior can be anticipated or changed. Simulation models must therefore portray decisionmaking behavior as it is, and not as it might be if decision-makers were omniscient optimizers. The decisionmaking heuristics and strategies people use, including their limitations and errors, must be modeled.

Bounded rationality and behavioral decision theory

An extensive body of theory and empirical data exists which documents the strategies and heuristics people use in a wide variety of decisionmaking contexts. The sources of this knowledge include organizational studies, cognitive and social psychology, and other social sciences. Known generally as Behavioral Decision Theory (BDT), these studies emphasize bounded rationality in human behavior. BDT focuses on identifying cognitive limitations in the perception and processing of information and the organizational strategies people devise to deal with these limitations (Armstrong 1985, Hogarth 1980, Kahneman et al. 1982, Simon 1982). BDT not only illuminates the way decisions are actually made, but documents a large number of systematic deviations from objectively rational behavior. Many heuristics lead to suboptimal or biased decisions in a wide variety of settings. Common examples include the gambler's fallacy and the regression fallacy (Tversky and Kahneman 1974). BDT shows that "people give more weight to data that they consider causally related to a target object..." (Hogarth 1980, 42-43, emphasis in original). However, people are poor judges of causality and correlation, and in controlled experiments systematically create mental models at variance with the known situation. Ironically, most people, including many professionals, consistently assert that their own performances are immune from such pitfalls, are reluctant to abandon their mental models, and selectively use hindsight to 'validate' their preconceptions.

BDT is useful in simulation for two reasons. First, bounded rationality provides theoretical foundations for behavior that deviates from objective rationality. Second, the empirical results of BDT research document the heuristics people actually use, providing a data base for model development. However, the empirical results of BDT are overwhelmingly micro-level. It is difficult to connect the results of BDT to the aggregate decision rules typically used in simulation models. Compare, for example, the verbal protocols and models described in Feigenbaum and Feldman 1963 or Ericsson and Simon 1984 with the typical continuous simulation model. Protocols for decisionmaking heuristics are usually given in the form of decision trees or other discrete, event-oriented procedures such as the TOTE unit (Miller, Galanter, and Pribram 1960). In contrast, a typical continuous simulation decision rule for inventory management in a manufacturing firm might be:[2]

$$DP_t = EO_t + (DINV_t - INV_t)/TCI$$

where

DP = Desired Production (units/time)
EO = Expected Orders (units/time)
TCI = Time to Correct Inventory (time)
INV = Inventory (units)
DINV = Desired Inventory (units)

The continuous rule may be used to describe aggregate behavior for a firm or industry. It is not intended as a literal statement of how production decisions are made. Rather it is deemed to be an acceptable simplification. The lumping of distributed but similar components, as in the aggregation of stocks of different product lines and firms into a single measure of inventory, is often cited as justification for assuming continuous decision rules (Forrester 1961, Ch. 11). Such aggregation is justified as necessary if a model is to remain small enough to be comprehensible, and thus useful. Yet the inevitability of aggregation does not mean such aggregation is appropriate. (For informative discussion of the connection between representations of feedback at the event level with continuous representations, see Richardson 1984). The production scheduling example, though simple, shows that methods are needed to bridge the gap between the micro knowledge of individual decisions and the macrobehavior of aggregate phenomena. Simulation has long been touted as one such method, but its acceptance has been limited by the inability to relate the micro data to aggregate decision rules. Direct experiment offers a useful method to bridge this gap.

Experimental economics

Direct experimental investigation of economic behavior has flowered over the past two decades (for surveys, see Smith 1982a, Plott 1982; also Smith 1979, 1982b). Most of these experiments concern what Smith 1982a calls microeconomic systems. These microworlds consist of an environment and an institutional structure. The environment includes the number of agents participating and their individual preferences, knowledge, and resource endowments. The institutional structure of the experiment consists of a specified language for interaction, resource allocation rules, cost imputation rules, and adjustment process rules governing the beginning, transitions, and end of the experiments. By manipulating both the environment and the institutional arrangements, the experimenter creates controlled situations in order to test hypotheses or elicit new data. Typical experiments investigate aspects of price theory such as the number of buyers and sellers required to find competitive equilibrium, test whether decisionmaking is conducted in accordance with expected utility theory, and evaluate the efficiency of various institutions such as different types of auctions. While many of these studies are concerned with equilibrium or asymptotic results, a relative few studies are primarily concerned with dynamic behavior (e.g. Plott and Wilde 1982, Garner 1982, Alker and Tanaka 1981, Williams 1979). As Shubik (1979, 354) notes:

> When an economy is in equilibrium, the role of markets, financial institutions and money tends to disappear. The institutions such as organized markets, firms and banks are the carriers of process and a major part of the information and communication flow of an economy. In disequilibrium they appear clearly....Many different institutions may have the same static efficiency properties, but it is possible that they manifest considerably different dynamic properties. The questions concerning the selection of optimal...institutions in a fully dynamic context have hardly been asked in a precise form, let alone answered.

Because experimental microeconomics must pay careful attention to the institutional structure of the system, that is, to the procedural aspects of decisionmaking, it is inherently dynamic and well suited to test the decision rules of behavioral simulation models, even if they are not strictly microeconomic in focus. The need for explicit specification and control of the information set available to the actors and the rules of interaction and exchange give the experimental method the flexibility required to test the behavior of real people against the models of behavior assumed in simulations.[3]

Protocols for direct experiment to test simulation models

The structure of models considered for experimental testing can be divided into two components: the physical and institutional structure on the one hand, and the behavioral decision rules on the other. For example, the physical and institutional assumptions in a model of a manufacturing firm might include the aggregation of different product lines into a single inventory of finished products from which orders are filled. Other assumptions may be that there is a certain average lag required to produce goods, that labor is the sole factor of production, that list prices are announced publicly. The behavioral decision rules would include procedures for determining production goals, workweek, hiring and layoffs, and changes in prices. The design of the experiment will depend on which type of assumption is to be tested. The discussion below focuses on tests of the decision rules. Tests of the physical and institutional assumptions are considered in the concluding section.

To test the decision rules of the model, the experimenter must ensure that the human subjects are placed in the same physical and institutional context assumed in the model. The purpose of such a test is to determine whether real people behave in the same way the model presumes them to behave, given the physical structure and other aspects of the organizational setting.

Likewise, the behavioral decision rules in simulation models presume that a certain information set is available at each decision point. The managers of the firm in the example above may not know the demand schedule of the customers, but only the past history of orders and prices. Behavioral simulation models, in keeping with the theory of bounded rationality, often presume that decisions are factored into subdecisions, and that the local decisionmaking units may have access to or choose to use less than the full set of available information (Morecroft 1983, 1985). Typically, decisions would emphasize locally available, relatively certain, and relatively new information over distant, uncertain, or dated information.

The information available to the human subject must be carefully controlled. Several designs are possible here, depending on the purpose of the experiment. One can deliberately restrict the information available to the human subjects to the set assumed to be actually used in the model, so as to see if the live agents process that information in the same way as presumed in the model. Alternatively, one may give the live subjects more information than is presumed to be used in the model, and test whether they utilize the same subset of the available information.

The decision rules in behavioral simulation models also impute preferences to the decisionmakers. These preferences may not take the form of explicit cost or utility functions, but may appear as a set of subgoals which the decisionmaker strives to satisfy. For example, the production model described above may assume that desired inventory is determined by a desired coverage ratio of expected orders. Implicit in this subgoal is a loss function which assumes costs arise from excess of deficient inventory coverage. The experiment must ensure that subjects are faced with the same underlying preferences and costs the modelbuilder presumes to be operating. For example, to motivate players to balance inventory and desired inventory, the experimenter might create a cost function which specifies losses for excess or deficient inventory, and reward players for minimizing their costs. The cost function approach is taken in the "Beer Distribution Game" (Sterman 1984).

If the experimenter can control for the institutional structure, information availability, and preferences assumed in the model, then the resulting behavior of live subjects can be directly compared to the behavior produced by the decision rule of the model. The human subjects are placed in the same physical and institutional structure, given the same information set, and strive for the same goals as the simulated decisionmakers. But whereas the simulation model contains an explicit rule for processing the information to yield a decision, the subjects of the experiment are free to make their decisions any way they wish. The comparison of simulated and experimental behavior thus provides a potential disconfirmation of the model's decision rules (Bell and Senge 1980). A strong caveat must be issued here, however. The correspondence of experimental and simulated behavior does _not_ validate the model--after all, any of the assumptions about physical structure, information availability, and preferences may be false. Additional experiments are necessary to test these assumptions. A successful outcome shows only that _given_ the institutional structure, real people behave the same way the model presumes them to behave.

An example: A behavioral model of capital investment cycles

The example presented here involves a model of aggregate investment behavior. The original simulation model (Sterman 1985) showed how the capital investment decisions of individual firms could lead to long-period cycles in the economy. A simple model of a capital-producing firm was developed. The decision rules of the model were shown to be locally rational through partial model tests. That is, the individual decision rules behaved rationally in isolation. The response of the partial model to unanticipated shocks was smooth, stable, and appropriate. Next, a macroeconomic linkage (the capital investment accelerator) was introduced. The accelerator represents the fact that capital is an input to its own production. When the demand for capital rises, capital-producing firms must expand their own capacity, further increasing the total demand for capital (cf Frisch 1933, Samuelson 1939, Goodwin 1951). Introducing the accelerator into the model caused large amplitude limit cycles to emerge (Figure 1). The model showed that locally rational decisionmaking by individual firms could lead to macroeconomic instabilities when the firms were coupled through the accelerator mechanism.

The physical and institutional structure assumed in the model is extremely simple (Figure 2). Orders for the firm's product accumulate in the backlog, which is depleted by production. Production is determined by capacity and capacity utilization. Utilization is a nonlinear function of the ratio of desired production to capacity: when desired output exceeds capacity, production is constrained by capacity; when desired output is less than capacity, utilization is gradually cut back. Capacity is determined by the capital stock and the (constant) capital/output ratio. Capital stock is augmented by acquisitions and diminished by discards. The average lifetime of capital is assumed to be constant and the discard process exponential. Orders for capital are received after a delay representing the construction process. Hence orders for capital accumulate in the supply line (the backlog of unfilled orders for capital, including units under construction). The supply line is diminished when construction is completed and the capital enters the productive stock.

The key decision rule in the model is the capital order decision. The firm must decide how much capital to order each time period given available information such as the current backlog, past order rates, capacity, capacity on order, and the capital acquisition delay. The assumed order decision is decomposed into several blocks of equations:

$$CO_t = C_t * COF_t \tag{1}$$

$$COF_t = f_1(ICOF_t) \qquad 0 \leq f_1 \leq COF_{max}, \qquad f_1' \geq 0 \tag{2}$$

$$ICOF_t = (CD_t + CC_t + CSL_t)/C_t \tag{3}$$

where

CO = capital order rate (capital units/year)
C = capital stock (units)
COF = capital order fraction (fraction/year)
ICOF= indicated capital order fraction (fraction/year)
CD = capital discard rate (capital units/year)
CC = correction to orders from capital stock (capital units/year)
CSL = correction to orders from supply line (capital units/year).

Three motivations for ordering capital are assumed: first, to replace discards; second, to correct any discrepancy between the desired and actual capital stock; and third, to correct any discrepancy between the desired and actual supply line. The sum of these three pressures, as a fraction of the existing capital stock, defines the indicated capital order fraction ICOF. The actual order fraction COF is a nonlinear function of the indicated order fraction. For indicated order fractions between 5%/year and 25%/year, COF=ICOF. In extreme circumstances, however, the indicated capital order fraction may take on unreasonable values. For example, an extreme excess of capacity could cause ICOF to be negative. But since gross investment must be positive, COF asymptotically approaches zero as ICOF drops below 5%/year. Similarly, to prevent the order fraction from taking on unreasonably large values, it is assumed that the maximum capital order fraction COF_{max} is 30%/year. The limit reflects physical constraints to rapid expansion such as labor and materials bottlenecks, financial constraints and organizational stress.

$$CSL_t = (DSL_t - SL_t)/TASL \qquad (4)$$

$$DSL_t = CD_t * PCAT_t \qquad (5)$$

$$PCAT_t = CAT_t \qquad (6)$$

where

DSL = desired supply line (capital units),
SL = supply line (capital units),
TASL = time to adjust supply line (years),
PCAT = perceived capital acquisition time (years),
CAT = capital acquisition time (years).

Firms strive to eliminate discrepancies between the desired and actual supply lines within the time to adjust supply line TASL. To ensure an appropriate acquisition rate, firms must maintain a supply line proportional to the delay they face in acquiring capital. If the acquisition time rises, firms must plan for and order new capital farther ahead, increasing the required supply line. For simplicity, the perceived capital acquisition time is assumed to equal the actual acquisition time.

$$CC_t = (DC_t - C_t)/TAC \qquad (7)$$

$$DC_t = RC * f_2(IC_t/RC), \quad f_2(0)=0, \quad f_2(1)=1, \quad f_2' \geq 0, \quad f_2'' \leq 0, \qquad (8)$$

$$IC_t = IPC_t * COR \qquad (9)$$

where

DC = desired capital (capital units),
TAC = time to adjust capital (years),
RC = reference capital (capital units),
IC = indicated capital (capital units),
IPC = indicated production capacity (units/year),
COR = capital/output ratio (years).

Like the supply line correction, firms attempt to correct discrepancies between desired and actual capital stock over a period of time given by the time to adjust capital. Desired capital is nonlinearly related to the indicated capital stock, which is the stock needed to provide the indicated production capacity IPC. Indicated production capacity is the capacity judged necessary to meet expected demand. Diminishing returns to capital are assumed to limit capital expansion when IC becomes large relative to the initial equilibrium capital stock RC.

$$IPC_t = EO_t + CB_t \tag{10}$$

$$CB_t = (B_t - IB_t)/TAB \tag{11}$$

$$IB_t = NND*EO_t \tag{12}$$

where

EO = expected orders (units/year),
CB = correction from backlog (units/year),
B = backlog (units),
IB = indicated backlog (units),
TAB = time to adjust backlog (years),
NDD = normal delivery delay (years).

Indicated production capacity reflects the capacity the sector judges necessary both to fill expected orders and adjust the backlog of unfilled orders to an appropriate level. The speed with which the sector strives to correct discrepancies between the actual and indicated backlog is determined by the time to adjust backlog. TAB represents management's sensitivity to abnormal delivery delays. Indicated backlog is the backlog required to supply the expected order rate within the normal delivery delay.

$$(d/dt)EO_t = (OR_t - EO_t)/TAO \tag{13}$$

where

EO = expected orders (units/year),
OR = order rate (units/year),
TAO = time to average orders (years).

The expected order rate represents the sector's forecast of demand. Adaptive expectations are assumed. Incoming orders are smoothed because it takes time for firms to decide that an unanticipated change in demand is lasting enough to warrant capacity expansion. Smoothing filters out short-term noise in demand, providing a more certain measure of long-run demand than the raw order rate and preventing wild swings in investment by allowing the backlog to buffer the system from the short-term variability of demand. First-order exponential smoothing is assumed for the averaging process. The smoothing time is given by the time to average orders TAO.

The assumed capital order decision represented by eq. 1-13 is typical of continuous simulation models. The rule deliberately abstracts from the discrete nature of individual decisions. The formulation is intended to capture the aggregate result of the investment decisions made by many firms. Orders are expressed as a continuous function of various inputs. Those inputs are restricted to information that is locally available to the decisionmakers (e.g. backlog, capacity). Information an individual firm is unable or unlikely to have, such as the value of the equilibrium capacity stock, is not used. The firm's forecasting process is rather simple. Finally, the formulation includes appropriate nonlinearities so that it is robust in extreme conditions: gross investment is constrained to be positive and finite; desired capital is assumed to reflect diminishing returns. The parameters assumed in the decision rule were chosen to be consistent with survey and econometric evidence reported in various studies (Mayer 1960; Coen 1975; Senge 1978, 1980).

Protocol for the experiment

To test the correspondence of the model to real behavior, an interactive simulation 'game' was developed. In the game, human subjects play the role of the manager of the capital producing sector and are responsible for making investment decisions. The structure of the game, the physical and institutional context in which the player makes decisions, is identical to the original model. The player is given the same information set available in the original simulation model. The only difference between the original simulation model and the game is the fact that investment decisions are specified in the latter by the player and in the former by the decision rule described above.[4]

The simulation game is described fully in Sterman and Meadows 1985. It can be played manually or on a personal computer. In the PC version, the game board is displayed on the screen, showing the current values of variables such as capacity, desired production, orders, etc. (Figure 3). The players enter their order decision for the current period in the box marked 'New Orders-Capital Sector'. Using animation, the flows of orders, shipments, and depreciation of the capital stock are graphically displayed on the screen.

The current values of all the system variables are displayed on the screen at all times. Players have the option of plotting and/or printing the entire history of the game to date at any time before entering their order decision. Thus perfect and complete information is available to the player. The only unknown in the system is the future order stream placed by the goods sector.

The player, or team of players, takes the role of manager for the entire capital-producing sector of the economy. Time is divided into two-year periods. At the beginning of each period, orders for capital are received from two sources: the goods sector and the capital sector itself. Orders for capital arriving from the goods sector are exogenous, as in the simulation model. Orders for capital from the capital sector are determined by the player. Orders placed by both sectors accumulate in the corresponding halves of the rectangle containing the supply line of unfilled orders. The sum of the supply lines of the two sectors is desired production for the current two year period. Production itself is the lesser of desired production or production capacity. Capacity is determined by the capital stock of the sector. Capital stock is decreased by depreciation and increased by shipments out of the supply line. Depreciation is 10% of capacity each period, corresponding to an average lifetime of 20 years. If capacity is inadequate to meet demand fully, available production of capital is allocated between the capital and goods sector in proportion to their respective backlogs. For example, if the backlog of the capital sector were 500 and the backlog of the goods sector were 1000, desired production would be 1500. If capacity were only 1200, production would be 1200 and the fraction of demand satisfied would be 1200/1500=80%. Thus 400 units would be shipped to the capital sector and 800 would be shipped to the goods sector. Any unfilled orders remain in their respective supply lines to be filled in future periods. In the example, 100 units would remain in the supply line of the capital sector and 200 would remain in the supply line of the goods sector.

Note that there is only one decision in the game that is left to the discretion of the player--how much new capital to order. The player's goal in making these decisions is to minimize the total score for the simulation. The score is the average absolute deviation between desired production and production capacity over the length of the game. Thus the score indicates how well the player has balanced supply and demand. Players are penalized equally for both excess demand and excess supply. The scoring rule supplies the loss function which shapes the preferences of the players.[5]

The games reported below were initialized in equilibrium with orders of 450 units each period from the goods sector and capital stock of 500 units. The capital sector must then order 50 units per period to replace depreciation. Desired production then equals 450 + 50 units, exactly equal to capacity, and yielding an initial score of zero.

Results

Several typical games are plotted in Figure 4; the results of 50 games are summarized in Table 1. The sample includes MIT undergraduate, master's, and doctoral students; PhD scientists and economists from various institutions in the US, Europe, and the Soviet Union; and business executives including several presidents and CEO's.[6] In all the games, orders for capital from the goods sector rise from 450 to 500 in year 4, and remain at 500 thereafter. The step change in orders is not announced to the players in advance.

Consider figure 4g. The player reacts aggressively to the increase in demand by ordering 150 units in year 4. The increase in orders further boosts desired production, leading the player to order still more. Because capacity is inadequate to meet the higher level of demand, unfilled orders accumulate in the backlog, boosting desired production to a peak of 1590 units in year 12, and slowing the growth of capacity. The fraction of demand satisfied drops to as low as 52%, so the player receives less than expected. Faced with high and rising demand, the player's orders reach 500 in the tenth year. Between years 14 and 16, capacity overtakes demand. Desired production then falls precipitously as the backlog of the capital sector is depleted, opening a large margin of excess capacity. Because of unfilled orders in the backlog, capacity continues to rise until year 18, reaching over 1600 units. Note that the step increase in goods sector orders raises total demand for capital by just 10 percent, but capacity reaches a peak more than 300 percent greater than its long-run equilibrium level. Faced with excess capacity, the player cuts orders back to zero. Capacity then declines through discards for the next 24 years. Significantly, the player allows capacity to undershoot its equilibrium value, initiating a second cycle of similar amplitude and duration. The other games shown in Figure 4 are much the same. The specifics vary, but the pattern of behavior in the games is remarkably similar. As shown in Table 1, the vast majority of players generate significant oscillations. The equilibrium value of capital stock is 560 units. The mean peak value of capacity was 2200, nearly 4 times the equilibrium level. The mean periodicity of the cycle is 45 years.

The oscillatory behavior seen in the majority of games is far from optimal. The optimal path (Figure 5) assumes that the shock is unanticipated: orders remain at their initial equilibrium level until <u>after</u> the rise in goods sector orders. Because capacity can only increase with a lag, the sudden increase in demand means the backlog of unfilled orders must rise above its equilibrium value. Thus capacity must rise above equilibrium to work off that excess backlog. After the backlog is reduced, capacity can fall back to the equilibrium value. In the optimal pattern, orders rise sharply immediately after the shock to boost capital stock above equilibrium and reduce the excess backlog. Unlike most actual games, the resulting rise in desired production does not cause further increases in orders. Instead, orders immediately drop below the replacement level, allowing capacity to fall back to the equilibrium level as the backlog of unfilled orders is filled. The optimal score is 19, thirty-one times less than the mean and 4.5 times less than the minimum score achieved in the sample of actual first-time players. Equilibrium is reestablished just 5 periods after the shock. In contrast, only 8 percent of the players were able to reach equilibrium within the 70 year time horizon of the game, even though there is no additional disturbance to the system after the initial shock, and it rapidly becomes clear that the goods sector will continue to order 500 units.

Comparing the simulation model with experimental results

To facilitate computation during the experiment, there are a few differences in parameters between the experiment and the original model (Note 4). To ensure that these differences do not produce spurious results the parameters of the model were altered to correspond exactly to the experiment. The behavior generated by the original simulation model, modified model, and the experiment is strikingly similar (Table 2, Figure 6). In all three cases,

1. Output rises slowly due to the lags in acquiring capital, but falls precipitously, followed by a long depression while the excess capital depreciates.
2. Capacity peaks after and higher than production.
3. Delivery delay peaks before the peak of output (the fraction of demand satisfied reaches its minimum before the peak of output).
4. Successive cycles occur despite the fact that there is no external disturbance after the initial step increase in orders.

The correspondence between the simulations and experimental results is excellent. The experiment shows that people do not behave rationally or optimally even when perfect knowledge of the system structure and perfect information are available, and even though the environment is highly simplified compared to real-life management situations. Indeed, players make basic errors such as ignoring the amount of capital on order, failing to anticipate the lag in acquiring capital, failing to realize that they will not receive everything they order within one period when the fraction of demand satisfied is less than one, and failing to anticipate the increase in the apparent demand caused by their own orders. Interestingly, few players exercised their option to plot out the behavior of the variables during the game.

The primary difference between the model and the experimental results is the fact that many of the players began to learn how to control the system as the game progressed. The mean capacity peak for the second cycle is "only"

1130 compared to 2200 for the first peak. The original model on the other hand presumes no learning, and generates a limit cycle that reaches a constant amplitude and persists without continuous exogenous triggering (Figure 1). It may seem that the diminishing amplitude of the cycles in the experiment reflects learning on the part of players. But this is not necessarily true. Note that the amplitude of the cycle in the game version of the model diminishes over time even though there is no learning process in the model (Figure 6). The parameters in the modified model result in a lightly damped cycle while those of the original model cause a limit cycle.[7]

Learning, however, does occur. It typically requires several plays of the basic game (with the single step increase in goods sector orders) for people to learn how to avoid the large amplitude cycles typical in their first play. And when the pattern of goods sector orders is changed, say by the inclusion of a small amount of random noise, the performance of the players deteriorates markedly; repeated play then brings the score down again. Players rapidly learn how to do better in the basic game, but an appreciation for the structure of the system and a robust ordering policy evolve more slowly.[8]

Is the experiment a fair test?

One might argue that the experiment, while perhaps interesting, does not reveal anything about investment behavior in the real world. The time available for play is too short, the problem too simplified. Further, real managers have access to decision aids such as corporate staffs, management information systems, and sophisticated models of the economy. This issue cannot be settled without further experimentation, but the parameters of the problem can be estimated (generalizing from experimental to field settings is discussed in Locke 1986). Investment decisions in the game are made in less time than is available for real investment decisions. But the decisionmaking task is much simpler than that for any actual investment, and the information available far more complete. These two effects offset each other. The issue is whether the time available for decisionmaking in the experiment is long enough with respect to the difficulty of the task. The experiment would not be very illuminating if the results were contingent upon speeding players through so fast that their short-term memory was overloaded, causing them to make errors they would not make if they had more time to reflect. A rough calculation suggests this not to be the case. First-time players typically complete the 36 periods in about 40-80 minutes, implying about 70-130 seconds per decision. There are eight pieces of information on the screen at any given time. Short-term memory can store 7 ± 2 chunks of information. It takes about 5-10 seconds to transfer a chunk between short- and long-term memory. 70-130 seconds seems to be adequate time to scan and store the data and manipulate it to produce an investment decision, even if subjects transfer several pieces of data to long-term memory, particularly since not all decisions take the same length of time. The first several require a long time, as participants familiarize themselves with the game. Decisions when capacity is inadequate also require relatively longer. But when capacity vastly exceeds demand, players make their order decisions (usually to order 0) quite rapidly, often in just a few seconds. Of particular significance here is the fact that the experimental protocol does not impose any overt time pressure on the participants--they proceed at their own pace.

<u>Why does the model work?</u>

Why does the model, with its highly aggregate, simplified representation of decisionmaking, correspond so well to the behavior of real players? Why does such a gross description of decisionmaking work at all? The task in the game is a member of the large family of stock management control problems. In such a problem, the decisionmaker strives to adjust some stock to a desired level, and to compensate for disturbances in the environment. Often there are lags in the response of the stock to control actions. In the game, the decisionmaker must adjust the level of capital toward some desired level and keep it there once the desired level is attained, taking into account the fact that capital depreciates and that there is a lag in acquiring new capital. The ordering rule in the model says simply "order enough to replace depreciation, modified by some fraction of the discrepancy between the desired and actual levels of capacity, and don't forget to take the supply line of previous orders into account." It includes obvious nonlinearities to prevent negative or infinite orders. Any heuristic for managing a stock must take these motivations into account or fail in an obviously irrational manner. A decision rule that failed to replace the expected loss from the stock would produce steady state error in which the actual quantity would always be insufficient. A rule that failed to compensate for discrepancies between the desired and actual values of the stock could not respond to a change in the desired level of the stock and would allow the stock to drift randomly in response to environmental disturbances. The replacement and stock adjustment motivations for ordering are essential. In addition, a rule that fails to adjust for the supply line of capital on order will always overorder, producing instability.[9]

The decision rule in the model works because it captures the essential attributes of any reasonable stock-management procedure. No matter how detailed or complex the actual decisionmaking procedure is, it must compensate for depreciation and adjust for discrepancies between the desired and actual stock. The excellent fit between the aggregate rule and the behavior of real people reflects what Simon (1969) calls the near decomposability of the system:

> ...We knew a great deal about the gross physical and chemical behavior of matter before we had a knowledge of molecules, a great deal about molecular chemistry before we had an atomic theory, and a great deal about atoms before we had any theory of elementary particles....
>
> This skyhook-skyscraper construction of science from the roof down to the yet unconstructed foundations was possible because the behavior of the system at each level depended on only a very approximate, simplified, abstracted characterization of the system at the level next beneath....
>
> Artificial systems and adaptive systems have properties that make them particularly susceptible to simulation via simplified models....Resemblance in behavior of systems without identity of the inner systems is particularly feasible if the aspects in which we are interested arise out of the organization of the parts, independently of all but a few properties of the individual components (Simon 1969, 17).

In other words, it is the feedback structure of the system that determines its behavior, not the details of the decision rules.

Caveats and conclusions

The experiment shows that the continuous, aggregate decision rules used in behavioral simulation models can be excellent representations of real behavior. But the results reported here do not validate the model. The validity of the model is contingent on assumptions about both the institutional structure and the decision rules used by actors in the system. The experiment shows the behavior of human subjects is not significantly different from the behavior of the decision rule for investment assumed in the model. But it leaves the institutional assumptions untested.

The experimental method described here can be used to test these institutional assumptions. Among the most important are the aggregation of all capital-producing firms into a single sector under the control of the player, and the perfect information thus made available. These assumptions can be tested by re-designing the game for multiple players. In such a design each player would order capital from a supplier and receive orders from various customers. The individual players would be linked by an input/output matrix specifying interfirm transactions. But as in real life, an individual player would not be able to distinguish final demand from orders caused by transient stock adjustments or acceleration effects. A successful outcome would help build confidence in the appropriateness and utility of the original model. Failure to replicate would show the aggregation assumptions of the original model to be flawed.

The experimental method described here thus provides a process for building confidence in models where primary data on decisionmaking behavior are unavailable and significant aggregation is inevitable. It offers a useful tool for reproducible testing of hypotheses about institutional structure and decisionmaking behavior. Further, the experimental method seems to offer a promising approach toward building confidence in simulation models at all levels. While the approach may be particularly useful in simulations of macro systems where more usual research methods are not feasible, it should be useful in models of smaller organizations as well.

NOTES

1. Nelson and Winter 1982 provide a notable exception based on bounded rationality.

2. See e.g. Holt et al. 1960, Forrester 1961, Mass 1975, and Lyneis 1980, for models employing this or similar rules for production scheduling.

3. Experimental economics is closely related to simulation gaming. The literature on participatory simulation games is large and diverse (for surveys, see Horn and Cleaves 1980 and Wolfe 1985). Gaming in system dynamics contexts includes the "Beer Game," a production-distribution game (Sterman 1984), and STRATEGEM-1, an economic development game (Meadows 1985). Unlike experimental economics, for the most part these games are designed as teaching aids for the education of the players and not as research tools.

4. There are minor differences in parameters between the two models to facilitate computation in the game.

Parameter	Original	Experiment
COR Capital/Output Ratio (years)	3	2
CAT Capital Acquisition Time (years)	1.5	2
DT Computation Interval (years)	small	2

In addition, all numbers are rounded in the experiment to the nearest 10 units. These differences substantially reduce the complexity of the player's decisionmaking task without influencing the essential dynamics of the game.

5. Average absolute deviation was used rather than quadratic or other possible loss functions solely for simplicity. The experiment could easily be replicated with alternative scoring rules to test robustness with respect to this assumption.

6. No monetary rewards were used to motivate the players, in violation of Smith's (1982a) protocol for experimental microeconomics. While the experiment can be replicated with financial incentives, players in the sample here reported that they took the game seriously and tried their best. Particularly for the academic and business players, pride and fear of embarrassment seemed to be strong motivators. Many players expressed chagrin at their performance; some attempted to destroy their first results and substitute later trials. It is important to debrief players and explain the causes of the instability to convert the frustration of playing into useful learning.

7. Sensitivity tests of the parameters are presented in Sterman 1985. The parameters most influential for stability are the capital/output ratio COR (eq. 9) and the stock adjustment parameters TAC and TAB (eq. 7 and 11).

8. The existence of significant learning in the real economy is open to question. The structure of the actual economy is far more complex, and information far less available than in the experiment. Interconnections among firms are not fully appreciated. Individual firms cannot distinguish, as a player in the game can, the 'true' long-run demand from the 'false' orders generated by transient stock-adjustments and self-ordering. Note that the optimal behavior in the game demands that investment fall dramatically just when demand is highest relative to supply. Do real firms scale investment back just when backlogs are bulging, demand growing, prices and profits high, and delivery schedules stretching out? In addition, the long time required in real life for the consequences of the accelerator to manifest reduces the likelihood that corporate and government managers will learn from experience. Learning is hindered by the low weight accorded to the record of past decades and the advice of 'elder statesmen' compared to the memory of recent events and the pressures of the moment. Note that the three-to-five games typically required to learn how to bring the system smoothly into equilibrium in the experiment corresponds to several hundred years of simulated personal experience in a controlled environment.

9. In fact many players do forget to take the supply line into account, exacerbating the instability (e.g. Figures 4a, 4b). The importance of the supply line correction is tested in Sterman 1985.

REFERENCES

Alker, H. and A. Tanaka (1981) "Resolutional Possibilities in 'Historical' Prisoner's Dilemmas," Working Paper, MIT Center for International Studies.

Armstrong, J. S. (1985) Long Range Forecasting. New York: Wiley.

Bell, J. and P. Senge (1980) "Methods for Enhancing Refutability in System Dynamics Modeling," TIMS Studies in the Management Sciences 14, 61-74.

Coen, R. (1975) "Investment Behavior, the Measurement of Depreciation, and Tax Policy," American Economic Review 65, 59-74.

Ericsson, K. A. and H. A. Simon (1984) Protocol Analysis: Verbal Reports as Data. Cambridge: The MIT Press.

Feigenbaum E. A. and J. Feldman (1963) Computers and Thought. New York: McGraw-Hill.

Forrester, J. W. (1961) Industrial Dynamics. Cambridge: The MIT Press.

Frisch, R. (1933) "Propagation Problems and Impulse Problems in Dynamic Economics," Gordon, R. and L. Klein (eds) Readings in Business Cycles. Homewood, Illinois: Irwin, 155-185.

Garner, C. A. (1982) "Experimental Evidence on the Rationality of Intuitive Forecasters," in Smith, ed. Research in Experimental Economics Volume 2.

Goodwin, R. (1951) "The Nonlinear Accelerator and the Persistence of Business Cycles," Econometrica 19, 1-17.

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

Hogarth, R. (1980) Judgement and Choice. New York: Wiley.

Horn, R. and A. Cleaves (1980) The Guide to Simulation/Games for Education and Training. Beverly Hills: Sage.

Kahneman, D., et al. (1982) Judgement under Uncertainty: Heuristics and Biases. Cambridge: Cambridge University Press.

Keynes, J. M. (1939) "Professor Tinbergen's Method," The Economic Journal. 49(195), 558-568.

Leamer, E. (1983) "Let's Take the Con out of Econometrics," American Economic Review 73 (1), 31-43.

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

Leontief, W. (1982) "Academic Economics," Science, 217 (9 July), 104-107.

Locke, E. (1986) Generalizing from Laboratory to Field Settings. Lexington, MA: DC Heath.

Lyneis, J. (1980) Corporate Planning and Policy Design. Cambridge: The MIT Press.

Mass, N. (1975) Economic Cycles: An Analysis of Underlying Causes. Cambridge: The MIT Press.

Mayer, T. (1960) "Plant and Equipment Lead Times," Journal of Business 33, 127-132.

Meadows, D. L. (1985) "STRATEGEM-1: A Resource Planning Game" Environmental Education Report and Newsletter, 14 (2), 9-13.

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

Morecroft, J. (1983) "System Dynamics: Portraying Bounded Rationality," Omega, 11 (2), 131-142.

Morecroft, J. (1985) "Rationality in the Analysis of Behavioral Simulation Models," Management Science. 31(7), 900-916.

Nelson, R. and S. Winter (1982) An Evolutionary Theory of Economic Change. Cambridge: Belknap Press of Harvard University Press.

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

Plott, C. R. (1982) "Industrial Organization Theory and Experimental Economics," Journal of Economic Literature 20, 1485-1527.

Plott, C. R. and L. Wilde (1982) "Professional Diagnosis versus Self-Diagnosis: An Experimental Examination of Some Special Features of Markets with Uncertainty," in Smith, ed. Research in Experimental Economics Volume 2. Greenwich: JAI Press.

Richardson, G. P. (1984) "The Evolution of the Feedback Concept in American Social Science," Ph.D. dissertation, Sloan School of Management, Massachusetts Institute of Technology.

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

Senge, P. M. (1978) "The System Dynamics National Model Investment Function: A Comparison to the Neoclassical Investment Function," Ph.D. dissertation, Sloan School of Management, Massachusetts Institute of Technology.

Senge, P. M. (1980) "A System Dynamics Approach to Investment Function Formulation and Testing," Socioeconomic Planning Sciences 14, 269-280.

Shubik, M. (1979) "On the Number of Types of Markets with Trade in Money: Theory and Possible Experimentation," in Smith, ed. Research in Experimental Economics Volume 1. Greenwich: JAI Press.

Simon, H. A. (1969) The Sciences of the Artificial. Cambridge: The MIT Press.

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

Simon, H. A. (1982) Models of Bounded Rationality Volume 2: Behavioral Economics and Business Organization. Cambridge: The MIT Press.

Smith, V. L. ed. (1979) Research in Experimental Economics Volume 1. Greenwich: JAI Press.

Smith, V. L. (1982a) "Microeconomic Systems as an Experimental Science," American Economic Review 72(5), 923-955.

Smith, V. L. (1982b) Research in Experimental Economics Volume 2. Greenwich: JAI Press.

Sterman, J. D. (1984) "Instructions for Running the Beer Distribution Game," D-3674, MIT System Dynamics Group.

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

Sterman, J. D. and D. Meadows (1985) "STRATEGEM-2: A Microcomputer Simulation Game of the Kondratiev Cycle," Simulation and Games 16(2), 174-202.

Thurow, L. (1983) Dangerous Currents. New York: Random House.

Tversky, A. and D. Kahneman (1974) "Judgement under Uncertainty: Heuristics and Biases," Science 185, 1124-1131.

Williams, A. W. (1979) "Intertemporal Competitive Equilibrium: On Further Experimental Results," in Smith, ed. Research in Experimental Economics Volume 1. Greenwich: JAI Press.

Wolfe, J. (1985) "The Teaching Effectiveness of Games in Collegiate Business Courses," Simulation and Games 16 (3), September, 251-288.

Figure 1. Limit cycle generated by original simulation model

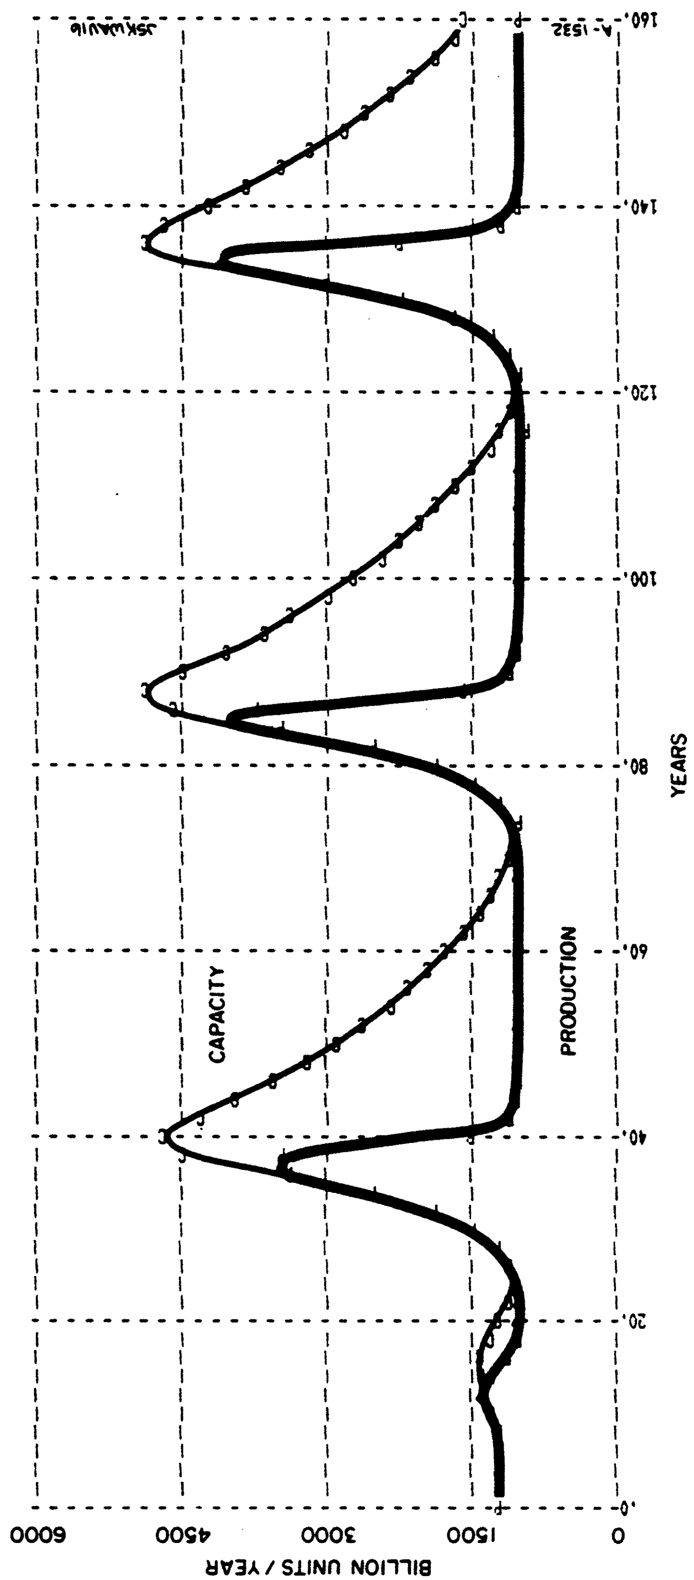

Figure 2. Structure of original simulation model.

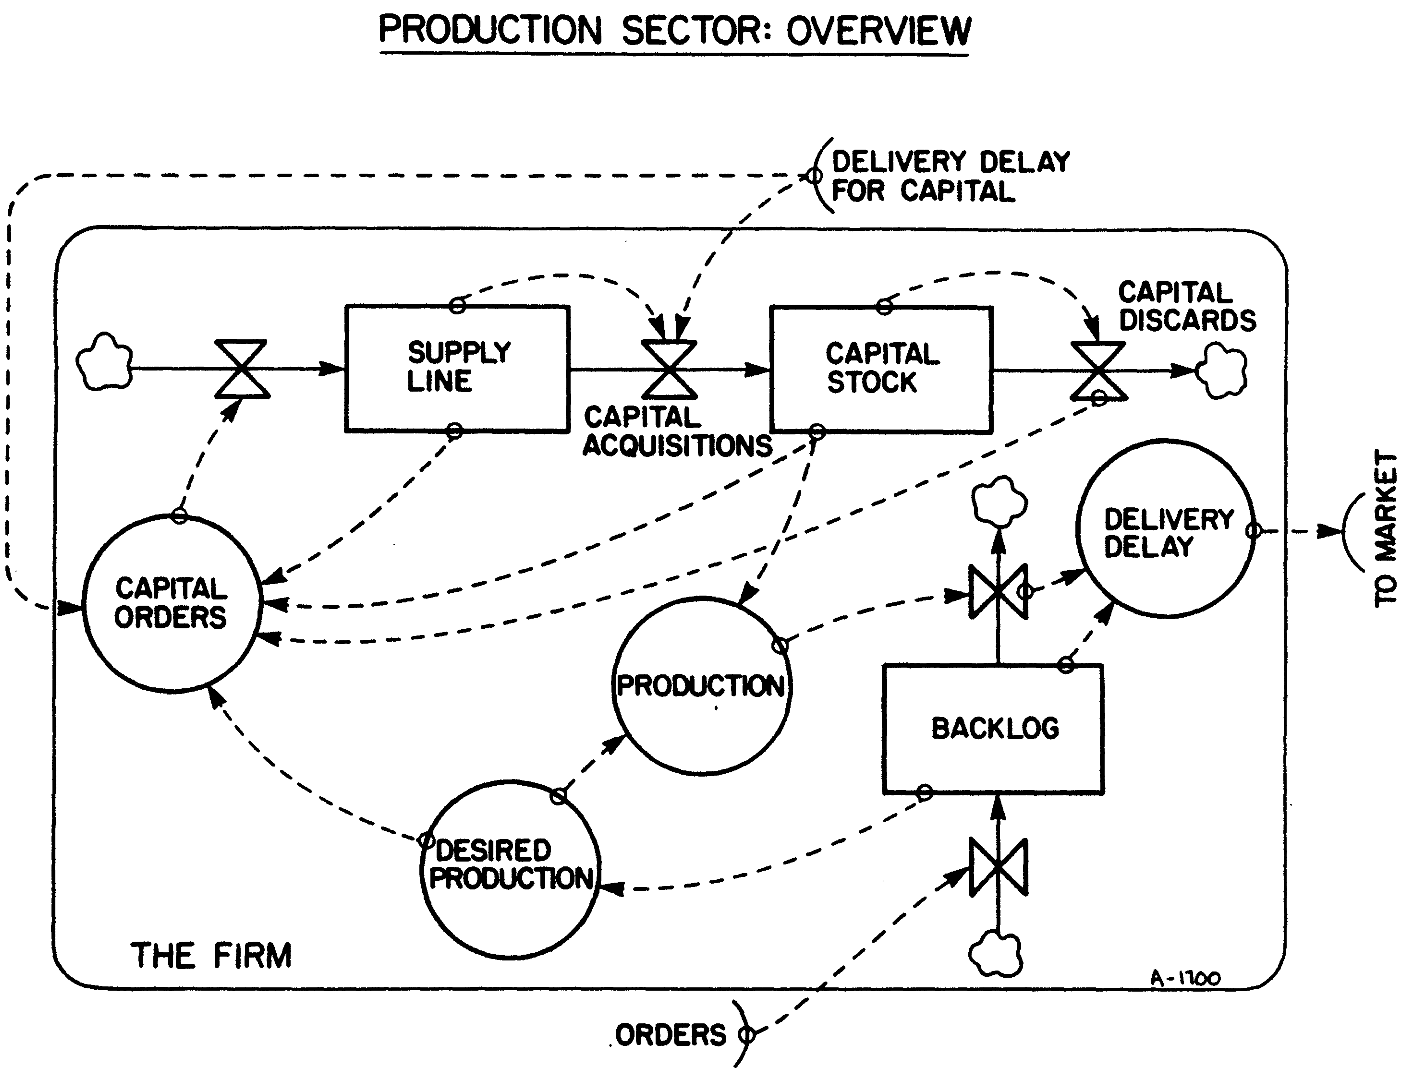

Figure 3. Game board, showing initial configuration.
Player is about to enter new orders for capital sector.

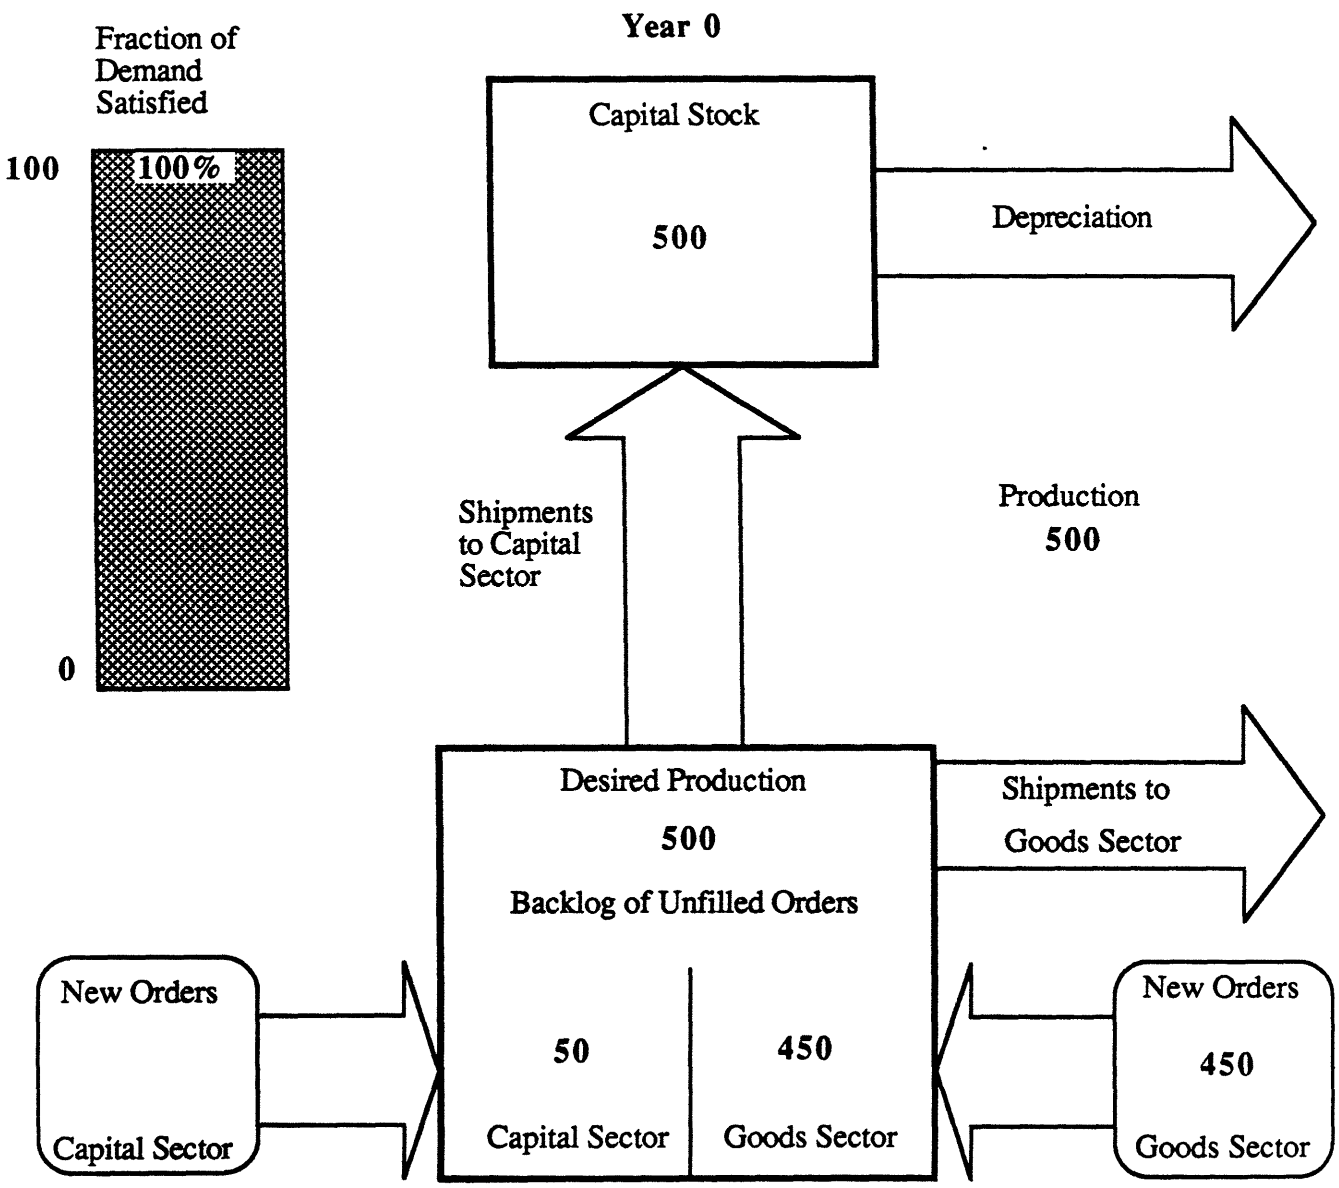

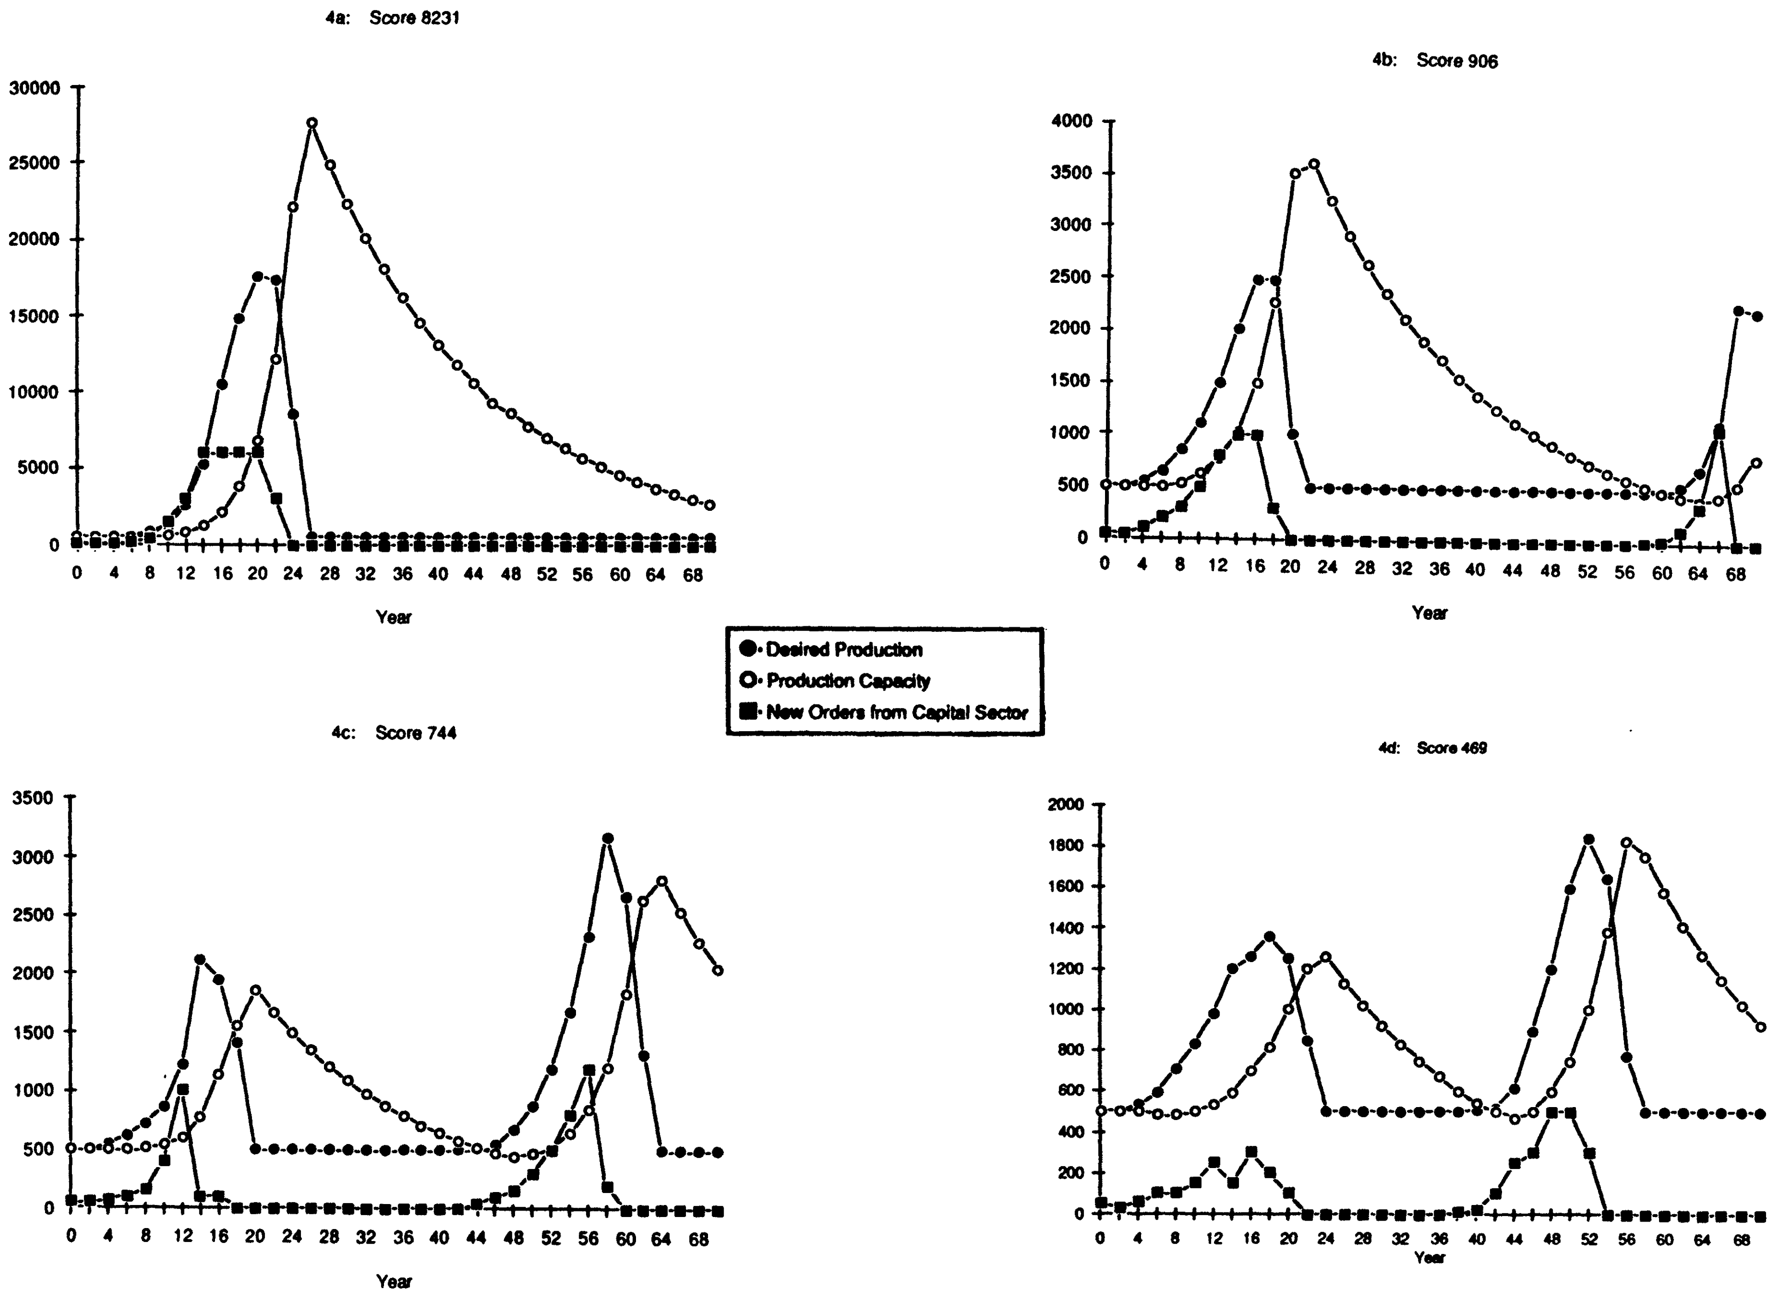

Figure 4. Experimental results. Note: vertical scales differ.

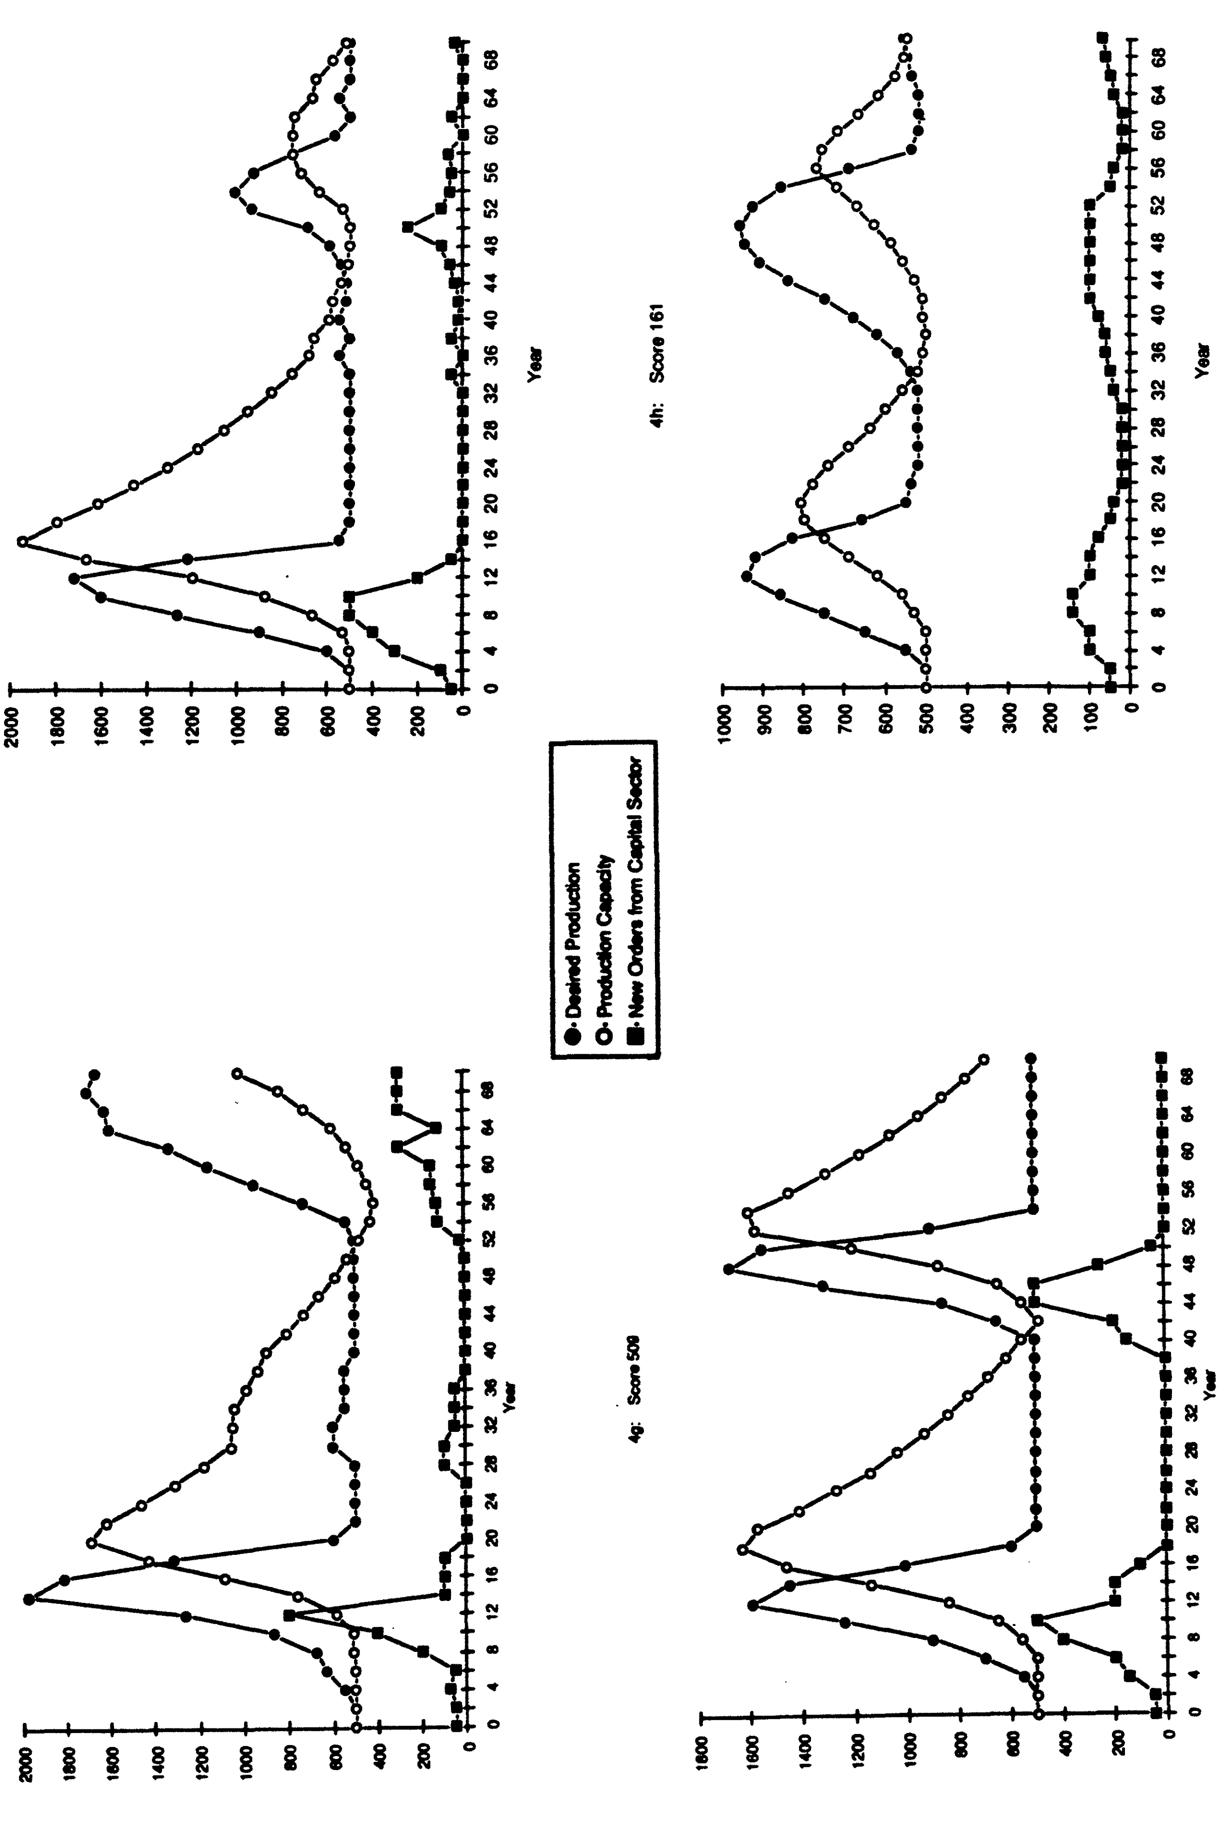

Figure 4 (cont.). Experimental Results. Note: vertical scales differ.

TABLE 1A
GAME RESULTS: 49 FIRST-TIME PLAYERS

	SCORE (UNITS/PERIOD)	PERIOD (YEARS)	1ST CAPACITY PEAK (UNITS/PERIOD)	2ND CAPACITY PEAK (UNITS/PERIOD)	PEAK ORDER RATE (UNITS/PERIOD)	MINIMUM ORDER (UNITS/PERIOD)	MINIMUM FRACTION OF DEMAND SATISFIED (%)
MEAN	597	46	2232	1139	629	4	48
MAXIMUM	8231	96[a]	27620	3020	6000	50	70
MINIMUM	86	26	600	560	90	0	20
STANDARD DEVIATION	1176	13	3935	671	927	11	14
N =	49	49	49	37[b]	49	49	49

TABLE 1B
GAME RESULTS: ABERRANT HIGH GAME EXCLUDED (SCORE 8231, FIGURE 4A)[c]

	SCORE (UNITS/PERIOD)	PERIOD (YEARS)	1ST CAPACITY PEAK (UNITS/PERIOD)	2ND CAPACITY PEAK (UNITS/PERIOD)	PEAK ORDER RATE (UNITS/PERIOD)	MINIMUM ORDER (UNITS/PERIOD)	MINIMUM FRACTION OF DEMAND SATISFIED (%)
MEAN	438	45	1703	1139	518	4	49
MAXIMUM	1992	84[a]	6390	3020	2200	50	70
MINIMUM	86	26	600	560	90	0	24
STANDARD DEVIATION	383	11	1346	671	501	11	13
N =	48	48	48	37[b]	48	48	48

---

a Periodicity was measured in production capacity. Periods greater than 70 were inferred from the game results assuming optimal play after year 70. This procedure is likely to introduce a downward bias in the reported period.

b 12 players generated so much excess capacity they had no opportunity to reach equilibrium or generate a second cycle before the end of the simulation.

c The game shown in Figure 4a represents a clear outlier. The score was 8231, orders reached 6000 units/period, and capacity reached a peak of 27,620, more than 49 times the equilibrium level.

Figure 5. Optimal behavior in the game. Note the vertical scale.

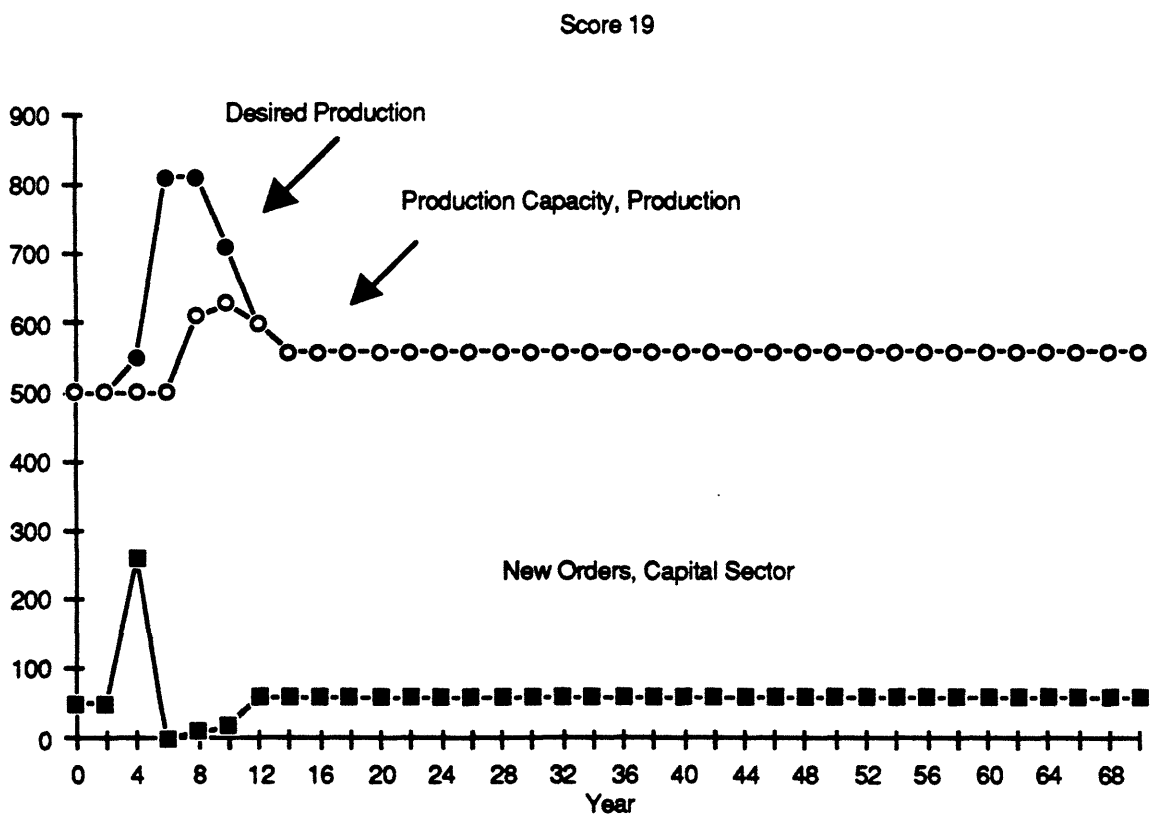

Figure 6. Behavior of modified model. Parameters of original model modified to correspond exactly to the experiment.

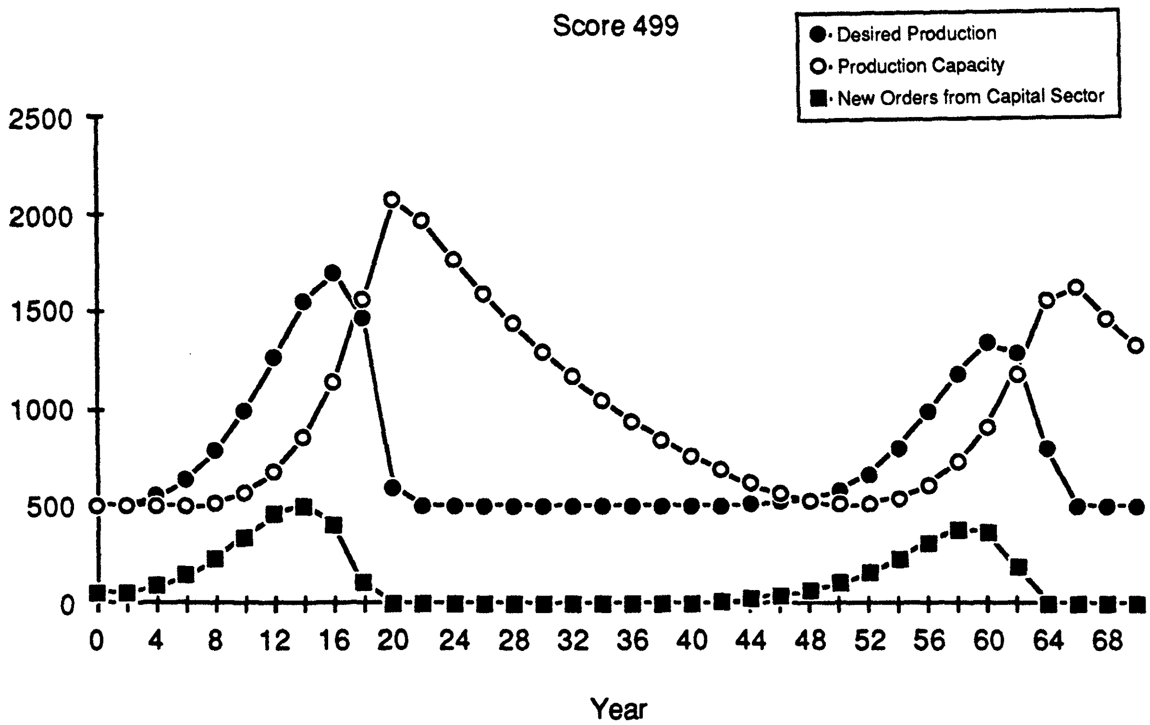

TABLE 2

COMPARISON OF ORIGINAL MODEL, MODIFIED MODEL, AND GAME RESULTS

	ORIGINAL MODEL	MODIFIED MODEL[a]	GAME RESULTS		GAME RESULTS, HIGH GAME EXCLUDED	
			MEAN	STD DEV	MEAN	STD DEV
PERIOD (YEARS)	49	46	46	13	45	11
SCORE (UNITS/PERIOD)	NA	499	597	1176	438	383
1ST CAPACITY PEAK (UNITS/PERIOD)	2318[b]	2072	2232	3935	1703	1346
2ND CAPACITY PEAK (UNITS/PERIOD)	2318[b]	1633	1139	671	1139	671
PEAK ORDERS (UNITS/PERIOD)	746[b]	502	629	927	518	501
MINIMUM ORDERS (UNITS/PERIOD)	0[b]	0	4	11	4	11
MINIMUM FRACTION OF DEMAND SATISFIED (%)	38[c]	53	48	14	49	13

a Computed from simulation shown in Figure 6.
b Computed from the simulation shown in Figure 1 by scaling the amplification generated in the original model to the equilibrium order rate in the game.
c Computed from the ratio of peak delivery delay to normal delivery delay in the original model.