D-3876-1

MISPERCEPTIONS OF FEEDBACK IN DYNAMIC DECISIONMAKING

by

John D. Sterman

WP-1899-87

D-3876-1

Misperceptions of Feedback in Dynamic Decisionmaking

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

May 1987

Presented at the 1986 meeting of the Judgment/Decisionmaking society, New Orleans, LA. The author wishes to thank Ernst Berndt, John Carroll, Robin Hogarth, and Don Kleinmuntz for valuable comments and assistance.

ABSTRACT

In recent years laboratory experiments have shed significant light on the behavior of economic agents in a variety of microeconomic and decision-theoretic contexts such as auction markets, portfolio choice, and preference elicitation. Despite the success of experimental techniques in the micro domain, there has been relatively little work linking the behavior of decisionmakers to the dynamics of larger organizations such as corporations, industries or the macroeconomy. This paper presents a laboratory experiment in which subjects manage a simulated economy. Subjects must invest sufficient capital plant and equipment to satisfy demand. Subjects were given complete and perfect information regarding the structure of the simulated economy, the values of all variables and the past history of the system. Nevertheless, the overwhelming majority of the subjects generate significant and costly oscillations. A simple decision rule based on the anchoring and adjustment heuristic is shown to simulate the subjects' decisions quite well. Several distinct sources of the subjects' poor performance are identified and termed "misperceptions of feedback." The decision rule is related to various models of economic fluctuations; implications for experimental investigation of dynamic decisionmaking in aggregate systems are explored.

In recent years laboratory experiments have shed significant light on human behavior in a variety of microeconomic and decision-theoretic contexts including auctions, bargaining, and preference elicitation (Plott 1986, Smith 1986, Slovic and Lichtenstein 1983, Grether and Plott 1979). Despite the success of experimental techniques in the domain of the individual and small group, there has been comparatively little work relating the behavior of decisionmakers to the dynamics of larger organizations such as an industry or the macroeconomy. Experiments in both economics and psychology have focussed (with significant exceptions, e.g. Hogarth and Makridakis 1981, Kleinmuntz 1985, Brehmer 1986, Smith 1986) on static and discrete judgments. Hogarth (1981, 198) emphasizes

> ...the continuous, adaptive nature of the judgmental processes used to cope with a complex, changing environment.... With few exceptions...judgment researchers have focussed on discrete incidents (particular actions, predictions, and choices) that punctuate these continuous processes; furthermore, task environments are typically conceptualized to be stable.... [I]nsufficient attention has been paid to the effect of feedback between organism and environment.

The complexity and scale of corporate and economic systems renders experiments on the systems itself infeasible. This paper argues that experimental studies of the "feedback between organism and environment" in aggregate dynamic systems such as the economy can be conducted in the laboratory with computer simulation models. Simulations can represent the structure and complexity of such systems with great fidelity and permit controlled manipulations of the decision context and information presented to the subject. A wealth of models exist which can provide an environment for the study of dynamic decisionmaking.

As an example, this paper presents a laboratory experiment in which subjects play the role of managers in a simple macroeconomic setting, specifically a model of aggregate capital investment. The subjects must maintain sufficient capital stock (plant and equipment) to satisfy demand. Subjects are given complete and perfect information regarding the structure of the simulated economy, the values of all variables and the past history of the system. Nevertheless, the behavior of the overwhelming majority of subjects is far from optimal. Most subjects generate significant, costly oscillations. A simple decision rule is proposed as a model of the heuristic used by the subjects. The rule is interpreted in terms of the anchoring and adjustment heuristic (Tversky and Kahneman 1974). The parameters of the rule are estimated for each trial using maximum likelihood methods, and the rule is shown to account for the subjects' decisions extremely well. The estimation results suggest an explanation for the poor performance of the subjects. The explanation takes the form of several distinct misperceptions of the feedback structure of the system. Finally, the decision rule is related to various theories of economic fluctuations, and implications for experimental investigation and simulation modeling of aggregate economic dynamics are explored.

Laboratory macroeconomics appears to be an oxymoron. Since the spatial and temporal scales of macroeconomic systems are too vast to permit experiments on the systems themselves, how can external validity be established in the face of the simplifications and compressions required to implement such experiments in the laboratory? The problems of external validity, of appropriate aggregation, of task design and so forth are the same in principle as those encountered in any research on judgment. Yet these difficulties in the exploration of dynamic macrosystems are perhaps more acute. These issues are raised in the concluding section.

The admitted difficulties in the experimental treatment of aggregate systems must be compared against the alternative. Macroeconomic models have long been criticized for ignoring primary data on decisionmaking behavior and for the use of ad hoc decision rules (Leontief 1971, Phelps-Brown 1972, Kaldor 1972). Empirical tests of macro models typically consist of econometric estimation using aggregate data. But the frequent failure of such tests to discriminate between radically divergent theories has bred great dissatisfaction among economists:

> "If the method cannot prove or disprove a qualitative theory, and if it cannot give a quantitative guide to the future, is it worthwhile?" (Keynes 1939, 566);

> "Year after year...the econometricians fit algebraic functions of all possible shapes to essentially the same sets of data without being able to advance, in any perceptible way, a systematic understanding of the structure and the operations of a real economic system" (Leontief 1982, 107);

see also Leamer (1983), Thurow (1983). The work presented here does not call for the rejection of econometrics. But the low power of aggregate statistical tests warrants exploration of a complementary approach: laboratory experiments with simulated macroeconomic systems.

Background: The multiplier-accelerator theory of investment

The macroeconomic system chosen for experimental investigation is the multiplier-accelerator model of capital investment. First treated formally by Frisch (1933) and Samuelson (1939), the model synthesizes two important concepts, one at the firm level, and one at the macro level. At the level of the individual firm, it has long been recognized that a firm's demand for capital (plant and equipment) fluctuates far more than the demand for the goods produced with that capital stock. The principle of acceleration explains why this must be so. Imagine a manufacturing firm in equilibrium. For simplicity (and without loss of generality) assume that capital is the only input to the production process. Now consider the response of the firm to an unanticipated, permanent increase in the demand for its product. In the long run, the firm must increase its production capacity, and hence its capital stock, in proportion to the demand increment.[1] To do so it must temporarily increase the rate of gross investment (orders for new plant and equipment) above the rate of discard of old capital (figure 1). Investment must exceed discards long enough to build the capital stock up to the new equilibrium level. Typically the amplitude of the change in investment exceeds that of the change in final demand. As the firm's capacity approaches the new equilibrium, investment declines back towards the replacement level. Thus the accelerator is an inevitable consequence of the physics of capital accumulation within an individual firm.

While the accelerator is a micro-level phenomenon, the multiplier arises from the interconnections among firms in the economy. To illustrate, consider the adjustment process of the aggregate capital-producing sector of the economy, that is all the firms which produce plant and equipment. Imagine the effects, beginning in equilibrium, of an unanticipated, one-time increase in the demand for capital deriving from producers of consumer goods and services. Each firm in the capital sector, responding to the surge in orders, will seek to expand production capacity and must place orders for additional capital to do so. Such orders further increase the total demand for capital, causing still more pressure for expansion and stimulating additional orders, multiplying the demand increase of the consumer sector in a positive feedback loop (figure 2). Taken as a whole the aggregate capital-producing sector of the economy orders and acquires capital from itself. Hence the multiplier is the macroeconomic consequence of the fact that capital is an input to its own production.[2] Intuitively, such a positive feedback will amplify the tendency of capital investment to overshoot the long run level, leading to excess capacity, retrenchment and a drop in investment. The self-reinforcing nature of the multiplier feedback destabilizes the demand for and production of capital; as a result interactions between the multiplier and accelerator have long been central to several major theories of business fluctuations (Frisch 1933, Samuelson 1939, Goodwin 1951; Zarnowitz 1985 surveys recent theories).

But while multiplier-accelerator (MA) models have been extensively studied and estimated, and the concepts are taught in nearly every undergraduate macroeconomics course, the decision rules by which individual firms order capital stock have not been tested experimentally.

The assumptions behind the MA theory are of two types: the physical and institutional assumptions which characterize the structure of the economy on the one hand, and the behavioral decision rules imputed to managers on the other. The physical and institutional assumptions (the fact that capital is an input to its own production, the existence of construction and delivery delays for plant and equipment, and the average life of capital) are not much disputed. In contrast, assumptions about the decision processes by which managers choose when and how much to invest are highly controversial. Traditional models such as those of Samuelson as well as their econometric descendents (reviewed by Jorgenson, Hunter, and Nadiri 1970) typically assume that individual firms first decide how much capital they require, based on expected demand and static profit maximization criteria. They then order a fraction of the gap between their desired and actual stock each period until the actual stock equals the desired stock, taking into account the replacement of depreciation. The simplest such 'stock adjustment' model assumes a constant fraction α of the gap is corrected each period:

$$I_t = \alpha * (K^*_t - K_t) + \delta K_t \quad (1)$$

where I is gross investment, K^* and K are the desired and actual capital stocks, and δ is the fractional depreciation rate.

Many subtleties have been added to the basic stock adjustment rule in eq. (1), but these preserve the essence of the difference-reduction approach. Such rules have been criticized as ad-hoc. Critical economists charge (correctly) that such decision rules are not based on the optimizing motives and rationality which are the hallmark of microeconomics and the static theory of general equilibrium. More recent theories (e.g. Berndt, Morrison, and Watkins 1981, Meese 1980) address these defects by adding 'costs of adjustment' to the traditional static cost function of the firm. In such models, the costs borne by the firm depend not just on the price and quantity of the inputs used, but also on the rate of change of those inputs. Firms with rational expectations will choose investment to maximize the expected present value of profits π, as illustrated in eq. (2) for the case where capital is again the only input (time subscripts have been dropped for clarity):

$$\text{choose } I = K' - \delta K \text{ to maximize } \pi = \int_0^{\infty} \{P^*Q(P) - C[Q(P),K,K',P_k]\}e^{-rt}\,dt \qquad (2)$$

where C[·] is the firm's cost function, Q is the quantity of the firm's product demanded at the prevailing market price P, P_k is the price of capital, K' = dK/dt, and r is the discount rate. In theory, all the variables and parameters may be stochastic, and there may be arbitrarily complex feedbacks among them. In practice, severe simplifying assumptions are made (e.g. competitive factor and product markets, quadratic adjustment costs, etc.). Even so, as Pindyck and Rotemberg (1983) comment, "Stochastic control problems of this sort are generally difficult, if not impossible to solve. This, of course, raises the question of whether rational expectations provides a realistic behavioral foundation for studying investment behavior...."

Specifically, these models posit rational, optimizing motives *and* the ability on the part of managers to formulate and solve an exceedingly complex dynamic optimization problem. Such ability is contingent on (i) knowledge of the cost function facing the firm; (ii) knowledge of all future contingencies (or at least their probability), or equivalently, knowledge of the structure of the economy from which contingencies may be deduced (the rational expectations hypothesis [Muth 1961]); (iii) the cognitive wherewithal to solve the resulting optimization problem; and (iv) the time to do so. Thus while the modern theories of investment solve the problem of ad hoc decision rules, they do so by invoking assumptions about the motives and cognitive capabilities of managers which are in direct conflict with a vast body of experimental work in behavioral decision theory, cognitive psychology, and administrative science. Herbert Simon (1979, 510) puts it bluntly:

> 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 experiment described here offers the opportunity to test the adequacy of the models in eq. (1) and (2) directly.

METHOD

The experiment is based on a simple simulation model of the investment accelerator (Sterman 1985). The model was chosen because it explicitly portrays the physical and institutional structure described in figure 2, particularly the stock and flow structure of capital accumulation and production within an individual firm and the multiplier feedbacks between firms in the capital-producing sector. The model represents the aggregate capital-producing sector of the economy. Orders for capital arrive from two sources: the consumer goods sector and the capital sector itself. These orders are produced and shipped after a construction delay, provided the capital sector has adequate capacity. Capacity can be augmented by ordering new capital (which is received after the construction delay) and is diminished by depreciation of old capital. In the original model, a formal decision rule determined orders for new capital, closing the feedback loops in the system. In the experiment the rule is replaced by the subjects who are free to make investment decisions any way they wish as they attempt to balance supply and demand. Details of the conversion of the original model to a form suitable for experimental use appear in Sterman 1987.

The experiment is implemented on IBM PC-type microcomputers. A 'game board' is displayed on the screen (figure 3). Color graphics and animation highlight the flows of orders, production, and shipments to increase the transparency of the structure.

Subjects play the role of manager for the entire capital-producing sector of the economy. Each time period represents two years and follows the same sequence of events:

1. At the beginning of each period orders for capital are received from two sources: the consumer goods sector and the capital sector itself. The contents of the boxes labelled "New Orders–Capital Sector" and "New Orders–Goods Sector" flow via animation into the corresponding backlogs of unfilled orders, where they add to any orders remaining from prior periods. The orders of the goods sector are exogenous (a fact known to the subjects). The orders of the capital sector are the subjects' decisions from the prior period.

2. Desired production for the current period is calculated and displayed on the screen. Desired production is the sum of the backlog of unfilled orders of the two sectors.

3. Production is calculated. Production is the lesser of desired production and capacity, and

is displayed on the right of the screen. Production capacity for the current period is equal to the capital stock; its value is displayed in the box marked "Capital Stock".

4. Newly produced capital is delivered to each sector. If production capacity is sufficient to satisfy demand, each sector receives the total quantity of capital it desires. For example, if the backlog of the capital sector were 500 units and that of the goods sector 1000 units, desired production would be 1500. If capacity were 1500 or greater production would be 1500. The full supply line of each sector would be shipped. 1000 units would flow through the pipe labelled "Shipments to Goods Sector" and 500 units would flow through the pipe labelled "Shipments to Capital Sector", augmenting the capital stock. But if production capacity is insufficient to satisfy demand only a fraction of the orders in the supply lines can be shipped. The fraction of demand satisfied (FDS) determines how much of each supply line is shipped. FDS is the ratio of production to desired production, displayed both numerically and in the position of the 'thermometer' in the upper left of the screen. If capacity were only 1200 units in the example above, production would be limited to 1200. FDS would be 1200/1500 = 80%. The capital sector would receive 400 units and the goods sector 800. The 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 in the supply line of the goods sector.

5. Depreciation is calculated. Depreciation causes production capacity to decline by 10% each period. These units flow via animation through the depreciation pipe and out of the system.

6. Finally, the subject decides how much new capital to order and enters it in the box labelled "New Orders -- Capital Sector". After doing so, time is advanced and the cycle repeats.

Note that there is only one decision left to the discretion of the subject -- 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:

$$S = (1/T) \sum_{t=0}^{T} | DP_t - PC_t | \qquad (3)$$

The score indicates how well subjects balance demand and supply. Subjects are penalized equally for both excess demand and excess supply. Absolute deviation rather than, say, quadratic or asymmetric losses were used solely for simplicity and transparency. The function of the score is analogous to that of the profit maximization criterion in eq. (2): it dictates the optimal pattern of investment decisions. Departures from the optimal score provide a simple metric for the "rationality" of the subjects' behavior.

Note that the current values of all system variables are displayed on the screen at all times. Subjects have the option of examining a graph showing the entire history of their trial to date before entering their order decisions. They may do so as frequently as they wish. Thus perfect and complete information is available to the subjects. The only unknown is the future stream of orders placed by the goods sector. A pre-trial briefing covered the concept of the multiplier, explanation of the game board, rules, and scoring function. Questions about the mechanics and rules were answered before and during each trial. No time limits were imposed.

Wherever possible in converting the simulation model for experimental use parameter values were altered to simplify the decisionmaking task of the subjects. The original model was formulated in continuous time and used realistic parameters derived from surveys and econometric studies. In the experiment subjects place orders for capital at discrete intervals. Each period represents two years, a round number which approximates the average capital acquisition lag in the United States (Mayer 1960). The capital discard fraction is 10%/period, corresponding to an average lifetime of 20 years (Coen 1975). In addition, all quantities are rounded to the nearest 10 units. The appendix describes the equations of the model.

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; PhD 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. All subjects were fluent in English. To attain this diversity the experiment was run in several sites over a period of months. There is no evidence to suggest the results were influenced by the timing or locale of the trials. A caveat: no monetary rewards were used to motivate the subjects, in violation of Smith's (1982) protocol for experimental microeconomics. However, experiments in preference reversal (Slovic and Lichtenstein 1983, Pommerehne et al. 1982, Reilly 1982, Grether and Plott 1979) suggest performance is not materially affected by reward levels. The experiment could easily be replicated with cash rewards, but the subjects reported that they took the game seriously and attempted to do their best. Particularly for the academics and executives, pride and fear of embarrassment seemed to be powerful motivators. Many subjects expressed chagrin at their performance; several attempted to destroy their first results and substitute a later trial.

RESULTS

The trials reported below were run for 36 periods. All were initialized in equilibrium with orders of 450 units each period from the goods sector and capital stock of 500 units. Capital discards are therefore 50 units per period, requiring the capital sector to order 50 units each period to compensate. Desired production then equals 450 + 50, exactly equal to capacity, and yielding an initial score of zero. Orders for capital from the goods sector, the only exogenous input to the system, remain constant at 450 for the first two periods to allow subjects to familiarize themselves with the mechanics of the experiment. In the third period orders placed by the goods sector rise from 450 to 500, and remain at 500 thereafter. The step input is not announced to the players in advance.

Before presenting the results it is useful to consider the optimal response to the disequilibrating step increase in demand. When the demand of the consumer goods sector rises to 500 units per period, equilibrium production capacity rises to 560 units. 500 of these units are needed to satisfy the demand of the consumer goods sector, and 60 are needed to replace discards. (10% of 560 becomes 60 since all quantities are rounded to the nearest 10 units).

The optimal response is shown in figure 4. The optimal response assumes, consistent with the information available to the subjects, that the step increase in demand is unanticipated. Capital sector orders therefore remain at their initial level until *after* the demand shock. To reach the new equilibrium stock of 560 the order rate must exceed depreciation during the adjustment period. However, because capacity can only increase with a lag, the sudden increase in demand means the backlog of unfilled orders must rise above its final equilibrium value. Production, and hence capacity, must therefore rise above equilibrium long enough to work off the excess backlog. After the backlog is reduced capacity can fall back to its equilibrium value. In the optimal response, orders for capital rise immediately after the demand shock to quickly boost capacity and prevent a large backlog of unfilled orders from building up. Peaking at 260 units, orders immediately fall to zero, allowing capacity to fall back to equilibrium as the backlog is reduced. Orders then gradually approach equilibrium from below. The optimal score for the 36 periods is 19. Equilibrium is reestablished just 5 periods after the shock. As dictated by the acceleration principle, investment overshoots its equilibrium, but there is no oscillation.

The subjects behave quite differently. Several typical games are plotted in figure 5; table 1 summarizes the sample. Trial 16 is typical. The subject reacts aggressively to the increase in demand by ordering 150 units in year 4. The increase in orders further boosts desired production, leading the subject 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. The fraction of demand satisfied drops to as low as 52%, slowing the growth of capacity and frustrating the subject's attempt to satisfy demand. Faced with high and rising demand, the subject's orders reach 500 in the tenth year. Between years 14 and 16 capacity overtakes demand. Desired production falls precipitously as the unfilled orders are finally produced and delivered. A huge margin of excess capacity opens up. The subject slashes orders after year 10, but too late. Orders placed previously continue to arrive, boosting capacity to a peak of over 1600 units. Orders drop to zero, and capacity then declines through discards for the next 24 years. Significantly, the subject allows capacity to undershoot its equilibrium value, initiating a second cycle of similar amplitude and duration. Note that the demand shock raises the total demand for capital by just 10%, but capacity rises over 300% at its peak.

The other trials are much the same. While specifics vary the pattern of behavior is remarkably similar. As shown in table 1, the vast majority of subjects generated significant oscillations, even though there are no external disturbances to the system whatsoever after the initial step in demand, and it rapidly becomes clear that the goods sector will continue to order 500 units. Only 4 subjects (8%) were able to reestablish equilibrium before the end of the simulation. The mean value of the first capacity peak is 2200 units, more than 350 percent greater than the peak of the optimal pattern. The scores range from 78 to more than 8000. The mean score is 31 times greater than the optimal score; even the lowest score is more than four times the optimal performance.

Modeling the behavior of the subjects

The qualitative similarity of the results suggests the subjects, though not behaving optimally, used heuristics with common features. The decision rule proposed here as a model of the subjects' behavior was used in the original simulation model upon which the experiment is based (Sterman 1985) and is a variant of rules long used in simulation models of corporate and economic systems (Holt et al. 1960, Forrester 1961, Mass 1975, Lyneis 1980). The rule determines orders for capital as a function of information locally available to an individual firm. Such information includes the current desired rate of production, the current capacity of the firm, the rate of capital discards, the supply line of orders for capacity which the firm has placed but not yet received, and the capital acquisition delay:

$$CO_t = f(DP_t, PC_t, CD_t, SL_t, CAD_t). \tag{4}$$

The rule can be decomposed into several components. First, the rule accounts for the obvious constraint that gross investment must be nonnegative. Thus, actual capital orders CO are determined by the indicated capital order rate ICO only if ICO≥0:

$$CO_t = MAX(0, ICO_t). \tag{5}$$

The indicated capital order rate consists of three terms, each representing a separate motivation for investment. To maintain the existing capital stock at its current value, the firm must order enough to replace capital discards CD. The firm is assumed to adjust orders above or below discards in response to two additional pressures. The adjustment for capital AC represents the response to discrepancies between the desired and actual capital stock. The adjustment for supply line ASL represents the response to the quantity of capital in the supply line, that is, capital which has been ordered but not yet received:

$$ICO_t = CD_t + AC_t + ASL_t. \tag{6}$$

Firms are assumed to adjust orders for capital above or below the discard rate in proportion to the gap between their desired capital stock DK and the actual stock. Desired capital stock is determined from the desired rate of production DP and the capital/output ratio κ:

$$AC_t = \alpha_k * (DK_t - K_t) \quad (7)$$

$$DK_t = \kappa * DP_t \quad (8)$$

The adjustment for capital stock creates a simple negative feedback loop. When desired production exceeds capacity orders for capital will rise above discards until the gap is closed. An excess of capital similarly causes orders to fall below replacement until the capital stock falls to meet the desired level. The adjustment parameter α_k determines the aggressiveness of the firm's response, and must be nonnegative.

The adjustment for the supply line is structurally analogous:

$$ASL_t = \alpha_{sl} * (DSL_t - SL_t) \quad (9)$$

$$DSL_t = CD_t * CAD_t \quad (10)$$

where DSL = the desired supply line. To ensure an appropriate rate of capital acquisition a firm must maintain a supply line proportional to the capital acquisition delay. If the acquisition delay rises, firms must plan for and order new capital farther ahead, increasing the desired supply line. The desired supply line is based on the capital discard rate – a quantity readily anticipated and subject to little uncertainty. To illustrate the logic of the supply line adjustment, imagine an increase in desired capital. Orders will rise due to the gap between desired and actual capital stock. The supply line will fill. If orders in the supply line were ignored ($\alpha_{sl} = 0$), the firm would place orders through the capital stock adjustment, promptly forget that these units had been ordered, and order them again. The supply line adjustment creates a second negative feedback loop which reduces orders for new capacity if the firm finds itself overcommitted to projects in the construction pipeline, and boosts orders if there are too few. It also compensates for changes in the construction delay, helping ensure the firm recieves the capital it requires to meet desired production. α_{sl} controls the sensitivity of the firm to units in the supply line.

The decision rule in equations (5-10) is intendedly quite simple. The rule deliberately abstracts from the discrete character of individual investment decisions in the real world, consistent with its original role in continuous simulation models of aggregate phenomena. It also expresses the corrections to the order rate as linear functions of the discrepancies between desired and actual quantities. Undoubtedly the ordering rules of firms are more complex, and other work (e.g. Senge

1980) considers a number of subtleties. It is useful to think of the linear equation for ICO as the first term of the Taylor series expansion of the more complex, underlying investment rule.

Note that the decision rule allows firms to determine their orders on the basis of information which is locally available to the firm itself. Information an individual firm is unlikely or unable to have, such as the value of the equilibrium capital stock or the solution to the dynamic programming problem for optimal investment, is not used. The firm's forecasting process is rather simple: capacity is built to meet current demand. Finally, the rule includes appropriate nonlinearity to ensure robust results: orders remain nonnegative even if there is a large surplus of capital.

The rule can be interpreted as an example of the anchoring and adjustment heuristic. The capital discard rate forms an easily anticipated and relatively stable anchor under the control of the firm itself. Replacing discards will keep the capital stock of the firm constant at its current level (assuming the capital acquisition delay remains constant). Adjustments are then made in response to the adequacy of the capital stock and supply line. No assumption is made that these adjustments are in any way optimal or that firms actually calculate the order rate as given in equations (5-10). Rather, pressures arising from the factory floor, from the backlog of unfilled orders and disgruntled customers, and from commitments to projects in the construction pipeline cause the firm to adjust its investment rate above or below the level which would maintain the status quo. For an individual in the experiment the interpretation is parallel: replacing discards to maintain the status quo is a natural anchor. Adjustments based on the adequacy of the capital stock and supply line are then made. Again, there is no presumption that subjects explicitly calculate the adjustments using the formulae in the equations. The adjustment parameters α_k and α_{sl} reflect the firm's or subject's response to disequilibrium: large values indicate an aggressive effort to bring capacity and the supply line in line with their desired levels; small values indicate a cautious approach, or a higher tolerance for discrepancies between desired and actual stocks.

For both the real firm and the subjects, the hypothesis that decisions are made via an anchoring and adjustment heuristic such as the proposed rule is motivated by the observation that the complexity of determining the optimal rate of investment overwhelms the abilities of the managers and the time available to make decisions.[3] As argued below, the anchoring and adjustment rule proposed here is not only simple to apply but is also locally or intendedly rational.

Estimation

Testing the decision rule requires estimation of the adjustment parameters α_k and α_{sl}. All other quantities required to compute orders are given by the experimental data. The values of desired production, capacity, capital discards, and the supply line of unfilled orders are displayed on the screen at all times. The capital acquisition delay, required to compute the desired supply line DSL, is readily determined from the fraction of demand satisfied:

$$CAD_t = 1/FDS_t \tag{11}$$

(if the firm receives each period only half of the orders it has placed it will take two periods to empty the supply line).

The model is nonlinear. To estimate the model an additive disturbance term is assumed

$$CO_t = MAX(0, ICO_t) + \varepsilon_t \tag{12}$$

$$\varepsilon_t \sim N(0,\sigma^2).$$

The errors ε are assumed to be independent, identical, and normally distributed. In this case, maximum likelihood estimates of the parameters may be found by minimizing the sum of the squared errors Σe_t^2. Estimates for each trial were found by grid search of the parameter space, subject to the constraints $\alpha_k \geq 0$ and $\alpha_{sl} \geq 0$. Independence and normality of the errors ε implies the estimated parameters of such nonlinear functions are consistent and asymptotically efficient, and the usual measures of significance such as the t test are asymptotically valid (Judge et al. 1980).

Estimates for the 49 trials together with standard errors are given in table 2. Note that the function $CO=f(\cdot)$ does not contain an estimated regression constant. Thus the correspondence of the estimated and actual capital orders, not just their variance around a common mean, provides an important measure of the model's explanatory power. Since the residuals will not in general satisfy

$$\Sigma e_t = 0 \tag{13}$$

(simulated and actual capital orders need not have a common mean) the conventional R^2 is not appropriate. The alternative

$$R^2 = 1 - \Sigma e_t^2 / \Sigma CO_t^2 \tag{14}$$

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.

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%. All but two of the estimated capital stock adjustment parameters are highly significant. The supply line adjustment parameter is significant in 22 trials.

To illustrate the performance of the rule, figure 6 shows two trials for which the rule works well. Note in trials 5 (R^2 =.99) and 38 (R^2 =.94) how the decision rule captures the timing and magnitude of the order peaks and also the subjects' failure to raise orders early enough to prevent a second cycle. Figure 7 shows two trials for which the decision rule is least satisfactory. In trial 2 (R^2=.39) the decision rule works quite well during the first cycle. In year 30, however, the subject suddenly increases orders, despite the fact that capacity exceeds demand by a huge margin. Apparently the subject attempted to reduce the score by raising demand to match capacity, rather than waiting for capacity to fall. By year 36 the subject realized that additional orders only worsened the capacity surplus and cut orders back to zero. The decision rule cannot generate such behavior. Trial 12 (R^2 =.33) shows evidence of substantial learning: having been caught with overcapacity in the first cycle, the subject orders aggressively in year 52, then drastically cuts orders in year 58, waiting patiently for capacity to rise. This suggests the subject is placing more weight on the supply line in the second cycle. To test this hypothesis each cycle in trial 12 was estimated separately. The results are consistent with the hypothesis of learning. R^2 for each cycle is .85 and .47, respectively, a substantial improvement. In the first cycle α_k=2.00 and α_{sl}=1.88, indicating that the subject responded very aggressively to the shortfall of capacity. In the second cycle, α_k=1.00 and α_{sl}=2.00, showing the subject reduced the aggressiveness of the capital stock adjustment by a factor of two while slightly increasing the weight accorded the supply line. These latter parameters are highly stabilizing compared to those of the first cycle, consistent with the smaller amplitude of the second cycle.[4]

As a further test of the decision rule the experimental scores were compared to the score produced by simulating the decision rule using the estimated parameters. If the decision rule were perfect, the simulated and experimental scores would be equal, and regressing the simulated scores on the experimental scores would produce a slope of unity (t-statistic in parentheses; trial 1 is excluded as an outlier – its score is 14 times greater than the mean):

$$\text{Experimental Score}_i = \underset{(9.4)}{1.06} * \text{Simulated Score}(\alpha_k, \alpha_{sl})_i \qquad i = 2,...49 \qquad (15)$$

$R^2 = .21$

The slope of the relationship is highly significant and virtually equal to unity, indicating good correspondence between the decision rule and the experiment.

DISCUSSION

Why does the decision rule explain the subjects' behavior so well? Given its simplicity, why does it work at all? The task in the experiment is a member of the large class of stock management problems. In such problems, the decisionmaker seeks to maintain some stock or system state at a target level or within an acceptable range. The decisionmaker must compensate for disturbances in the environment. Often there are losses from the stock and lags in the response of the stock to control actions. Examples include managing inventories and cash balances in a corporation, regulating the temperature of a house or industrial process, guiding a car along a highway, controlling interest rates, and finding the right pace of presentation in a lecture.

The decision rule works because it captures the essential attributes of any reasonable stock management heuristic. A rule which failed to replace losses would produce a steady state error in which the stock would always be insufficient. Heuristics which failed to compensate for discrepancies between the desired and actual stock could not respond to a change in the target; the stock would follow a random walk as shocks bombard the system. The rule also accounts for the lag in the response of the stock to control actions (though surprisingly, many people apparently do not, causing instability).

Why did people use the decision rule instead of optimizing? Despite the gross simplifications of the model compared to real life, despite perfect information and knowledge of the structure of the simulated economy, the optimal path is at once too difficult to compute and too different from intuitive notions of reasonable strategy (it is difficult to stop ordering when the gap between demand and capacity is largest – figure 4). Optimal stock management requires a different strategy in each situation, since optimal behavior is a whole system property which depends crucially on the nature of the feedbacks among the system components. In contrast the proposed rule can be readily applied in a variety of stock management situations and vastly reduces the information, knowledge of system structure, and computational ability required.

Intended rationality of the decision rule

Simplicity alone does not explain why people use the heuristic embodied in the proposed decision rule. After all, the performance of most subjects is quite unstable and far from optimal. If instability is intrinsic to the rule it is difficult to argue that it reflects intendedly rational behavior or that it would survive in people's repertoire of judgmental heuristics. Simulation experiments can be used to test for the intended rationality of the rule (Morecroft 1985). Figure 8 shows two computer simulations of the decision rule. In both simulations the adjustment parameters α_k and α_{sl} are .55 and .40, respectively, the mean values of the estimated parameters. Figure 8a shows the full model as used in the experiment. The behavior is quite similar to that produced by the subjects. The large overshoot of capacity, successive cycles, periodicity, and score are all characteristic of the experimental results. In figure 8b the multiplier feedback has been cut. In consequence desired production is completely exogenous and the capital acquisition delay is constant. The test can be interpreted as the situation of an individual firm too small to influence the demand for its product or the availability of capital from its suppliers. Here the response to a 10% step increase in demand is stable, there are no oscillations, and equilibrium is reestablished rapidly. While not optimal the response is excellent. The results demonstrate the intended rationality of the decision rule. The decision rule does not recognize the existence of any feedbacks from the capital order decision to the demand for or availability of capital. When the environment is as simple as the decisionmaker presumes it to be the response of the system to shocks is reasonable and appropriate.

Misperceptions of feedback

If the decision rule is locally rational, the explanation for the poor performance of the subjects must be sought in the interactions between the decision rule and the feedback structure of the simulated economy. Close analysis of the experimental results and simulations reveals several distinct sources of poor performance. These are termed 'misperceptions of feedback' because they reflect a failure on the part of the decisionmaker to adequately assess the nature and significance of the causal structure of the system, particularly the linkages between their decisions and the environment.

1. *Misperception of time delays:*

Failure to appreciate time delays is reflected in two distinct facets of the experimental results. First, there is a strong tendency for subjects to be overly aggressive in their attempts to correct discrepancies between the desired and actual capital stock (that is, α_k is too large). Second, there is a strong tendency to ignore the time lag between the initiation of a control action and its full effect (that is, α_{sl} is too small).

Global stability analysis of the model (Sterman 1985, Rasmussen, Mosekilde, and Sterman 1985, Szymkat and Mosekilde 1986) confirms the strong effect of the capital stock and supply line adjustment parameters on the stability of the system. More aggressive response to capital stock discrepancies has a strong destabilizing effect; more aggressive supply line control is stabilizing. Intuitively, the more new capital ordered in response to a given capital stock shortfall (the larger α_k), the bigger the supply line will become before the capital stock rises to the desired level, and the greater the subsequent overshoot of capital stock will be as those orders are delivered. The positive feedback of the multiplier amplifies the destabilizing effects of aggressive capital stock adjustment: large orders further boost desired production, encouraging subjects to order still more. To the extent the supply line is considered (the larger α_{sl}) the capital order rate will be cut back as the supply line fills, preventing overordering.

To test the above argument about stability the estimated parameters were regressed on the log of the score. The score is a rough measure of instability: high scores indicate large gaps between desired production and capacity, indicating greater disequilibrium:

$$\ln(\text{Score}_i) = \underset{(42.8)}{5.3} + \underset{(5.1)}{1.7}*(\alpha_k)_i - \underset{(-3.7)}{1.1}*(\alpha_{sl})_i \qquad i=2,...49 \qquad (16)$$

R^2 = .43 F = 16.8

The results are highly significant and consistent with the formal analysis of the model: subjects with more aggressive capital stock adjustments and less aggressive supply line adjustments tended to have substantially higher scores. In light of the strong role of the supply line adjustment on stability, it is remarkable that the estimated supply line adjustment parameter is zero or not significant in fully 27 of the 49 trials, indicating that the majority of the subjects failed to take the supply line into account at all. It is as if the subject purchases a car, but has to wait for delivery. The next day, the garage still empty, the subject goes to the dealer and orders another one.

To illustrate, the estimated parameters suggest that subject 5 (figure 6) reacted very aggressively to the gap between desired production and capacity, and considered the supply line to a much smaller degree. As a result, the subject's initial orders further widened the gap between demand and capacity, leading to still greater orders. By putting so little emphasis on the supply line compared to the capital stock adjustment, the subject substantially overorders. In year 18, with capacity nearly equal to demand, the subject has an additional 2000 units in the supply line. When these are delivered, capacity jumps to a peak well above the peak of desired production and more than seven times the equilibrium level. The resulting score of 965, fifth highest in the sample, is 51 times greater than optimal. Subject 38, in contrast, is much more conservative, ordering only about a fifth of the gap between desired production and capacity each period. As a result, the multiplier feedback is weakened, demand does not climb so high, and the overshoot of capacity is reduced significantly. Subject 38 appears to ignore the supply line altogether, but the conservative ordering policy prevents too much excess capacity from developing, and produces a score nearly six times less than that of subject 5.

2. *Misperception of feedbacks from decisions to the environment:*

Figure 8 shows that the average parameters would produce excellent results if demand were exogenous. But demand is not exogenous. The multiplier feedback causes the environment to react endogenously to the decisions of the subjects. Their decision process, however, appears to be predicated on an exogenous environment. Thus many subjects were surprised that they did not receive all the capital they ordered as they tried to boost capacity. They were confused by the fact that placing orders to increase capacity seemed to worsen the gap between demand and supply. And they were further shocked that desired production suddenly dropped just when they thought they had finally caught up (figure 5). These phenomena are direct consequences of the multiplier loop, that is, the feedbacks from the subject's actions to the environment. In the long run, ordering more capital does increase capacity, but in the short run it adds to the total demand, worsening the shortfall. Ordering more capital also raises desired production further above capacity, reducing the fraction of demand satisfied and delaying delivery. During the period of inadequate capacity unfilled orders accumulate in the backlog, swelling desired production. When capacity finally overtakes desired production, these accumulated orders are shipped, and desired production falls.

Failure to appreciate the reflexive character of capital orders also explains one of the more remarkable aspects of the subjects' performance: the failure to prevent a second cycle by allowing capacity to undershoot its equilibrium value. Consider trial 5. Between years 20 and 56 there is tremendous excess capacity. The subject orders zero to reduce capacity as quickly as possible. Demand consists entirely of the 500 units requested by the goods sector. By year 58 capacity has fallen to 570, and the impending discard rate is 60 units. Anticipating the one-period lag in acquiring capital, the subject orders 60 units. If demand remained at 500, capacity would stabilize just above demand, and the subject would have achieved a low-score equilibrium. By ordering enough to offset discards, however, total demand rises to 560 just as capacity falls to 510. Capacity has suddenly become inadequate, initiating the second cycle. The subject was apparently adjusting capacity to meet current demand, and failed to realize that in equilibrium capacity must be sufficient to meet the demand of the goods sector *and* replace discards. Thus the subject aims for a target which is too low. The decision rule generates the same mistake. The desired capital stock is based on current demand and the desired supply line is based on current discards. In consequence, during the period of excess capacity the decision rule aims for a capacity target which is too low and fails to increase orders until it is too late, just as the majority of the subjects do. The decision rule initiates a second cycle because it does not consider the global equilibrium state or the feedbacks from the order decision to the demand for capital.

Do such misperceptions of feedback exist in the real world, or are they artifacts of the unfamiliar task of the experiment? This is largely an empirical question. Yet several points bear consideration. To avoid such misperceptions it is not sufficient that the information required to perceive the feedback structure of the system be available. In the experiment, complete knowledge of the structure is given, and there is no uncertainty regarding the values of the variables. It is also necessary to infer from that structure the dynamic characteristics of the system. This is equivalent to solving the system of dynamic equations. The system in the experiment, simplified though it is, constitutes a third-order nonlinear differential equation. The intuitive solution of such a system is, to say the least, difficult. Realistic systems are substantially more complex, more nonlinear, and involve more time delays. Thus it is reasonable to hypothesize that the misperceptions of feedback identified in the experiment would likely be exacerbated in the vastly more complex feedback environment of real systems.

There are numerous examples of stock management situations in which the supply line is ignored or unknown, leading to instability. A teenager's first experiences with alcohol are paradigmatic. Inexperienced drinkers, unaware of the time delay between taking a drink and its effect, frequently overshoot the acceptable level of intoxication. If the time frame for the dynamics is short, learning can be expected to dampen the instability over time. For most people experience gradually produces an appreciation for the "supply line" of alcohol which has been consumed but not yet had its effect, for the number of drinks required to reach a given state of intoxication, and for the decay rate. The result is diminution of the aggressiveness with which the discrepancy between the actual and desired state of drunkenness is approached (smaller α_k and larger α_{sl}). But here the feedback between decisions and results is swift, the nature of the supply line and the effects of alcohol are reasonably apparent, experience can be accumulated rapidly and is highly salient (particularly the morning after). These conditions are frequently not met in economic settings. In many situations the supply line is distributed among large numbers of competitors and is thus unknown to each individual firm, or the time required for learning exceeds the tenure of individual decisionmakers. Instability in such situations is chronic. The business cycle, the recurrence of speculative bubbles (Kindleberger 1978), and cycles of boom and bust in commodities, agriculture, and real estate (Meadows 1970, Hoyt 1933) provide ready examples. Experiments with simulation models of these systems would shed significant light on the generality and limitations of the proposed heuristic.

There is an analogy to Hardin's (1968) "tragedy of the commons" here. For any individual firm in a competitive economy, the environment may appropriately be viewed as exogenous. Yet the interactions among these individual firms create strong feedbacks, feedbacks which cause locally rational decisionmaking procedures to produce results which are not only unintended but globally dysfunctional. Of course, unintended behavior arising from systemic feedbacks is not new, nor must it be dysfunctional for society. Adam Smith's invisible hand is a negative feedback loop which leads each individual "to promote an end which was no part of his intention." Experimental investigation coupled with simulation analysis can help distinguish conditions in which unintended and undesirable behavior arises.

CONCLUSIONS

The results of this work have several implications for research in dynamic decisionmaking and economics. The results show it is possible to test experimentally the decision rules of corporate and economic models. A decision environment representing an important macroeconomic situation was created. Two conflicting theories of decisionmaking behavior in that situation exist in economics. Traditional models assume individual firms follow a difference-reduction heuristic for investment. Modern theories assume firms behave so that their behavior is optimal with respect to some intertemporal objective function. The experimental results strongly suggest that subjects do not behave optimally even when provided with perfect information and knowledge of the system structure. The results are explained well by a simple heuristic which assumes individual firms follow the difference-reduction strategy. Further, the results reveal several misperceptions of feedback: many subjects fail to adequately account for the delay between a control action and its effect, and fail to understand the feedback between their own decisions and the environment. The "open-loop" character of their decisionmaking exacerbates instability.

This work is exploratory, and it is necessary to consider its limitations. One objection is that the simulation model is too simple. All capital producers are aggregated into a single sector. It makes no allowance for order cancellations or for a variable lifetime of capital. It ignores prices, financial markets, and labor markets. There is no possibility for government policies to stimulate demand. These criticisms are all correct. It does not necessarily follow, however, that subjects would behave better in a more realistic model or that the inclusion of additional market forces would ensure optimal behavior. Indeed, one might argue that the vastly more complex feedback structure of more realistic models may further degrade subjects' performance. The method outlined here provides the possibility of directly testing the role of market forces in ameliorating or amplifying misperceptions of feedback in dynamic environments.

The aggregation assumptions, for example, may be tested by running the experiment with multiple players. Each subject would manage an individual firm. Subjects would be linked by an input/output matrix specifying inter-firm transactions. Performance as a function of the availability of information and extent of collusion among subjects can then be readily assessed. These extensions may confirm or disconfirm the robustness of the decision rule with respect to information availability and industry structure. In like fashion the robustness of the rule with respect to the complexity of the task or objective function may be tested.

One could argue that through learning or natural selection optimal strategies would evolve. This is an empirical question which can and should be explored experimentally. Preliminary results suggest that three to five trials with the simple step in consumer demand are required for most subjects to learn to avoid instability and produce a low score. But when these same subjects are then given a random, cyclical, or growing pattern of demand, their performance degrades substantially. Apparently it is difficult for experience alone to overcome misperceptions of feedback structure and produce a robust ordering heuristic. Study of the processes by which people form and revise mental models of the feedback structure of their environment would appear to be a fruitful extension of the present research.

Finally, it appears that the experimental exploration of dynamic decisionmaking strategies in aggregate systems is feasible. The fidelity and flexibility of simulation models enables the investigator to construct rich, complex decisionmaking environments. The results can be directly compared to formal models of behavior. Simulation and formal analysis can be used to test for the intended rationality of such models, can establish stability conditions, and can guide normative policy design. The marriage of experimental research on judgment with realistic simulation models thus offers a reproducible procedure to explore the endogenous generation of macrobehavior from the microstructure of complex systems.

APPENDIX: Equations of the simulated economy

The equations of the model are given below in the difference equation form used in the experiment. Computer disks suitable for the IBM PC and compatibles are available from the author.

$$PR_t = MIN(DP_t, PC_t). \quad (A1)$$

Production of capital for the current period PR is the lesser of desired production DP and production capacity PC. Desired production for the current period is determined by the backlog of unfilled orders B. The backlog accumulates new orders O less production, implicitly defining a normal delivery delay of one period:

$$DP_t = B_t \quad (A2)$$

$$B_{t+1} = B_t + (O_t - PR_t). \quad (A3)$$

Note that the actual delay between placing and receiving an order depends on the availability of adequate capacity so that PR=DP. If PR<DP, the unfilled orders remain in the backlog to be produced in future periods.

Production capacity is determined by the capital stock of the sector K and its productivity, given by the capital/output ratio κ:

$$PC_t = K_t / \kappa. \quad (A4)$$

To simplify the decisionmaking task of the subjects κ is set to one period. Thus production capacity for the current period simply equals the capital stock.

Capital stock is the accumulation of acquisitions of new capital CA less discards of old capital CD:

$$K_{t+1} = K_t + (CA_t - CD_t). \quad (A5)$$

Discards are proportional to the stock of capital. The average lifetime of capital is given by the reciprocal of the discard fraction δ:

$$CD_t = \delta * K_t. \quad (A6)$$

The discard fraction is 10% per period, corresponding to an average life of capital of 20 years (Coen 1975, Senge 1980).

Capital acquisition is determined by the sector's supply line (the backlog of unfilled orders for capital). The supply line of unfilled orders is augmented as new orders for capital are placed CO, and diminished when those orders are acquired:

$$SL_{t+1} = Sl_t + (CO_t - CA_t). \quad \text{(A7)}$$

The capital acquisition rate depends on the supply line and the average delay in the construction and delivery of those unfilled orders. To simplify the experiment, the normal capital acquisition time is set to one period. However, if the suppliers of capital have insufficient capacity to fully satisfy demand, acquisition of capital will be reduced according to the fraction of demand satisfied FDS:

$$CA_t = SL_t * FDS_t. \quad \text{(A8)}$$

When FDS<1, the firm receives a fraction of the capital it desires; the unfilled orders remain in the supply line to be delivered in future periods.

Thus the model represents the structure of capital investment in a typical firm. The firm receives orders from customers and seeks to maintain an adequate stock of capital to fill those orders, taking into account the fact that capital wears out and that new capital can only be acquired with a substantial lag. In the experiment, the structure is used to portray the aggregate capital producing sector of the economy. Incorporating this macroeconomic linkage creates several interdependencies.

1. The total demand for capital arises from two sources: the capital ordered by the consumer goods sector GCO and the capital ordered by the capital sector itself:

$$O_t = GCO_t + CO_t. \quad \text{(A9)}$$

The total backlog of the capital sector B is then the sum of the backlogs of unfilled orders of the consumer and capital sectors:

$$B_t = GSL_t + SL_t, \quad \text{(A3')}$$

where GSL is the supply line of unfilled orders for capital placed by the consumer goods sector.

2. The normal delay in acquiring capital is determined by the normal delivery delay of the capital sector itself; both are equal to one period.

3. The fraction of demand satisfied FDS is determined by the capacity of the capital sector relative to total demand:

$$FDS_t = PR_t / DP_t. \quad \text{(A10)}$$

When $PC \geq DP$, FDS=1. But when capacity is inadequate, production is constrained by capacity and shipments of capital to each sector will be reduced in proportion to FDS.

The acquisition of capital by the consumer goods sector is analogous to the capital sector:

$$GCA_t = GSL_t * FDS_t, \tag{A11}$$

$$GSL_{t+1} = GSL_t + (GCO_t - GCA_t), \tag{A12}$$

where GSL is the supply line of unfilled orders for capital placed by the consumer goods sector, and GCA is the rate at which such orders are filled. The formulation for FDS ensures that capital production and total acquisitions are always equal: PR = GCA + CA.

Note that available production of capital is allocated between the two sectors in proportion to their supply lines, implying equal access to the pool of available production. Capital producers do not know the extent to which their customers are coupled through inter-firm transactions, and certainly do not consult an input-output table to assign priorities on the basis of technical coefficients. The assumption of proportional capital allocation is appropriate for an approximately competitive economy.

NOTES

1. Assuming constant returns to scale and no changes in factor prices. Relaxing these assumptions does not alter the essential character of the adjustment dynamics.
2. The investment multiplier is distinct from the Keynesian consumption multiplier in which an increase in consumer demand stimulates employment and boosts income, further increasing demand as consumers spend some fraction of their pay increase.
3. Of course, the parallel argument for the anchoring and adjustment rule is not intended to gloss the differences in the two decisionmaking contexts. Unlike the experiment, real firms may take months to study investment opportunities and may use decision aids such as consultants or econometric models. On the other hand, their choice is vastly more complex than that in the experiment: real firms must choose among a nearly infinite variety of plant and equipment; must find, build relationships with, and select vendors whose products best meet the firm's quality, cost, functionality, reliability, and other criteria; they must discuss and make tradeoffs among these criteria; must arrange financing and consider tax implications; must coordinate the investment with human and other resources; in short they must make a host of decisions, down to the color of the paint, which subjects in the experiment do not (see Cyert, Simon, and Trow 1956). The hypothesis that the proposed anchoring and adjustment ruled applies in both cases does not require equivalence of the decisionmaking tasks but only the weaker condition that in both cases the determination of optimal choices exceeds the abilities of the decisionmakers.
4. The t-statistics for α_k and α_{sl} are, respectively, 28.4 and 21.9 in the first cycle, and 8.3 and 7.7 in the second. A caveat: the nonlinear estimator used here is efficient only asymptotically. Estimating each cycle separately reduces the number of datapoints from 36 to 18. While the estimation results are sensible, the small sample properties of the estimator are not known. More detailed analysis of learning should consider repeated trials, and is left for subsequent research.

REFERENCES

Berndt, E., C. Morrison, and C. Watkins (1981). "Dynamic Models of Energy Demand: An Assessment and Comparison," in E. Berndt and B. Field (eds) Measuring and Modeling Natural Resource Substitution. Cambridge: The MIT Press.

Brehmer, B. (1986). "Feedback in Dynamic Decisionmaking," presented at the 1986 annual meeting of the Judgment/Decisionmaking Society, New Orleans, November.

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

Cyert, R., H. Simon, and D. Trow (1956). "Observation of a Business Decision," Journal of Business. 29, 237-248.

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

Frisch, R. (1933). "Propagation Problems and Impulse Problems in Dynamic Economics," in R. Gordon and L. Klein, eds., Readings in Business Cycles. Homewood, IL: Irwin, 155-85.

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

Grether, D. and C. Plott (1979). "Economic Theory of Choice and the Preference Reversal Phenomenon," American Economic Review. 69(4), 623-638.

Hardin, G. (1968). "The Tragedy of the Commons," Science. 162 (13 Dec), 1243-48.

Hogarth, R. (1981). "Beyond Discrete Biases: Functional and Dysfunctional Aspects of Judgmental Heuristics," Psychological Bulletin. 90, 197-217.

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

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

Hoyt, H. (1933). One Hundred Years of Land Values in Chicago. Chicago: University of Chicago Press.

Jorgenson, D. , J. Hunter, and M. Nadiri (1970). "A Comparison of Alternative Econometric Models of Quarterly Investment Behavior," Econometrica. 38, 187-212.

Judge, G. et al. (1980). The Theory and Practice of Econometrics. New York: John Wiley.

Kaldor, N. (1972). "The Irrelevance of Equilibrium Economics," The Economic Journal. 82, 1237-55.

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

Kindleberger, C. (1978). Manias, Panics, and Crashes. New York: Basic Books.

Kleinmuntz, D. (1985). "Cognitive Heuristics and Feedback in a Dynamic Decision Environment," Management Science. 31(6), 680-702.

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

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

Leontief, W. (1982). "Academic Economics," Science. 217 (9 July), 104-7.

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

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

Mayer, T. (1960). "Plant and Equipment Lead Times," Journal of Business. 33, 127-32.

Meadows, D.L. (1970). Dynamics of Commodity Production Cycles. Cambridge: MIT Press.

Meese, R. (1980). "Dynamic Factor Demand Schedules for Labor and Capital under Rational Expectations," Journal of Econometrics, 14, 141-58.

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

Muth, J.F. (1961). "Rational Expectations and the Theory of Price Movements," Econometrica, 29, 315-335.

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

Pindyck, R. and J. Rotemberg (1983). "Dynamic Factor Demands and the Effects of Energy Price Shocks," American Economic Review. 73(5), 1066-79.

Plott, C. (1986). Laboratory Experiments in Economics: The Implications of Posted Price Institutions," Science. 232 (9 May), 732-38.

Pommerehne, W. (1982). "Economic Theory of Choice and the Preference Reversal Phenomenon: A Reexamination," American Economic Review. 72, 569-75.

Rasmussen, S., E. Mosekilde, and J. Sterman (1985). "Bifurcations and Chaotic Behavior in a Simple Model of the Economic Long Wave," System Dynamics Review. 1(1), 92-110.

Reilly, R. (1982). "Preference Reversal: Further Evidence and Some Suggested Modifications of Experimental Design," American Economic Review. 72, 576-84.

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

Senge, P. (1980). "A System Dynamics Approach to Investment Function Formulation and Testing," Socioeconomic Planning Sciences," 14, 269-80.

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

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

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

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

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

Sterman, J.D. (1987). "Testing Behavioral Simulation Models by Direct Experiment," Management Science. Forthcoming.

Szymkat, M. and E. Mosekilde (1986). "Global Bifurcation Analysis of an Economic Long Wave Model," presented at the 2nd European Simulation Congress, Antwerp, Belgium, Sept. 1986.

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

Tversky, A. and D. Kahneman (1974). "Judgement Under Uncertainty: Heuristics and Biases," Science. 185, 1124-31.

Zarnowitz, V. (1985). "Recent Work on Business Cycles in Historical Perspective: A Review of Theories and Evidence," Journal of Economic Literature. 23, 523-80.

Figure 1. The Accelerator Principle. The simulation shows the typical response of a firm to an increase in its desired capital stock. For the capital stock to rise smoothly to the desired level, investment must overshoot the rate of capital discards. Note the amplification of investment relative to the change in desired capital: a 20% increase in desired capital produces a peak increase in investment of about 100%, yet the full adjustment still requires about 5 years (faster adjustment would cause more amplification). Such amplification is typical and explains the volatility of investment relative to final demand in the macroeconomy.

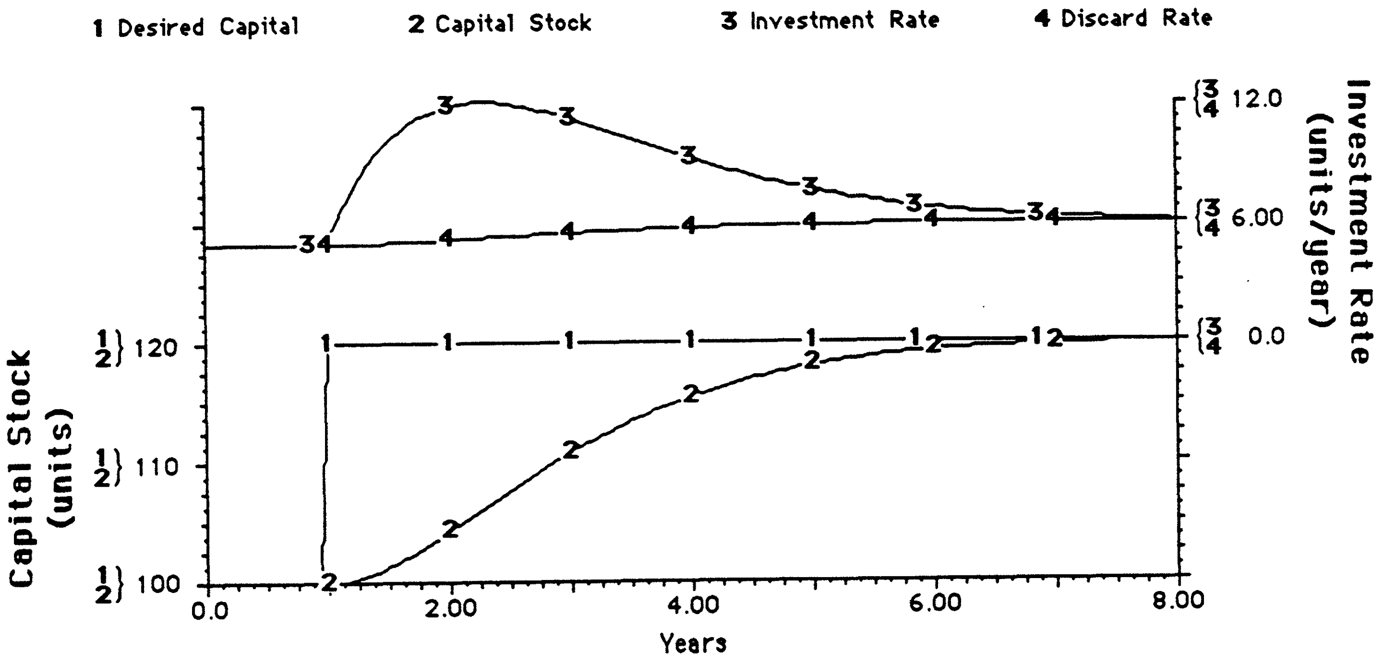

Figure 2. The Capital Investment Multiplier. Producers of plant and equipment order and acquire capital from one another. An increase in autonomous orders for capital causes a self-reinforcing increase in the total demand as firms in the capital-producing sector attempt to increase their production capacity. The demand for capital is multiplied by the fact that capital is required to produce capital.

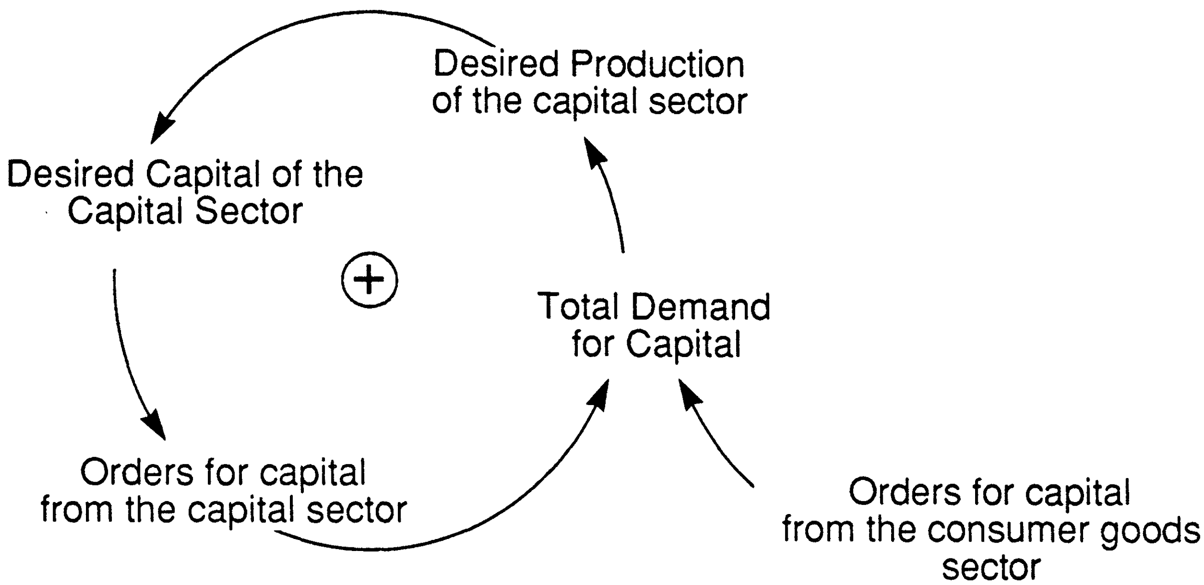

Figure 3. Computer screen as presented to the subject, showing the initial configuration. The subject is about to enter new orders for the capital sector. Color graphics and animation highlight the flows of orders, shipments, and depreciation.

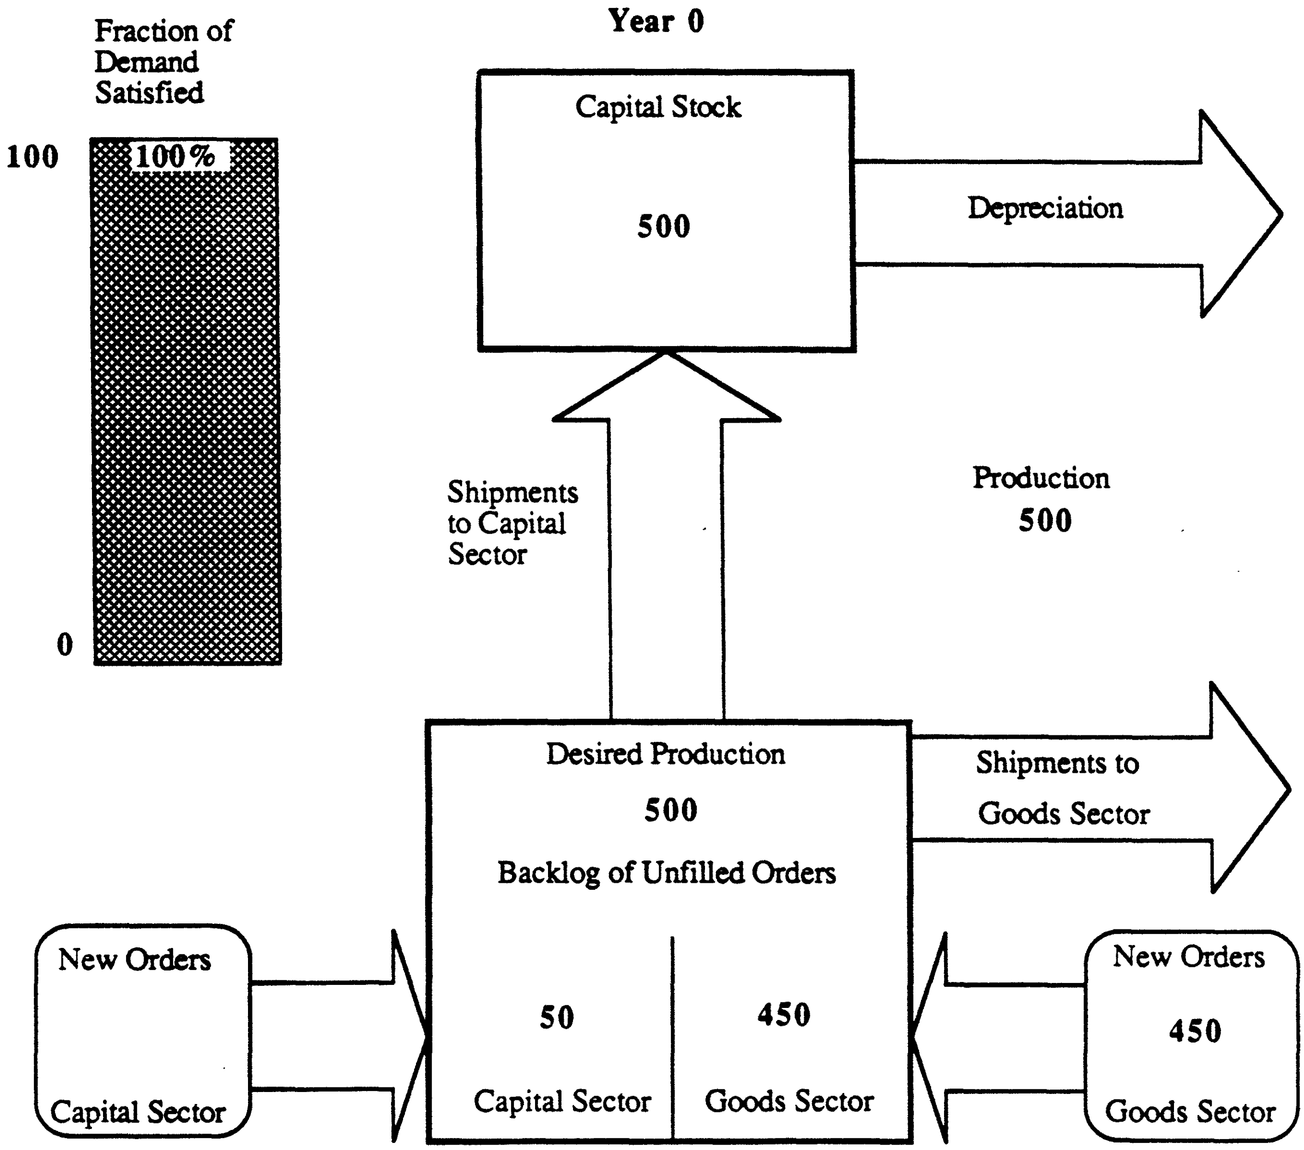

Figure 4. Each trial begins in equilibrium. In year 4 there is an unannounced increase in new orders placed by the consumer goods sector from 450 to 500 units per period (top graph). The optimal response (bottom graph) returns the system to equilibrium by year 14.

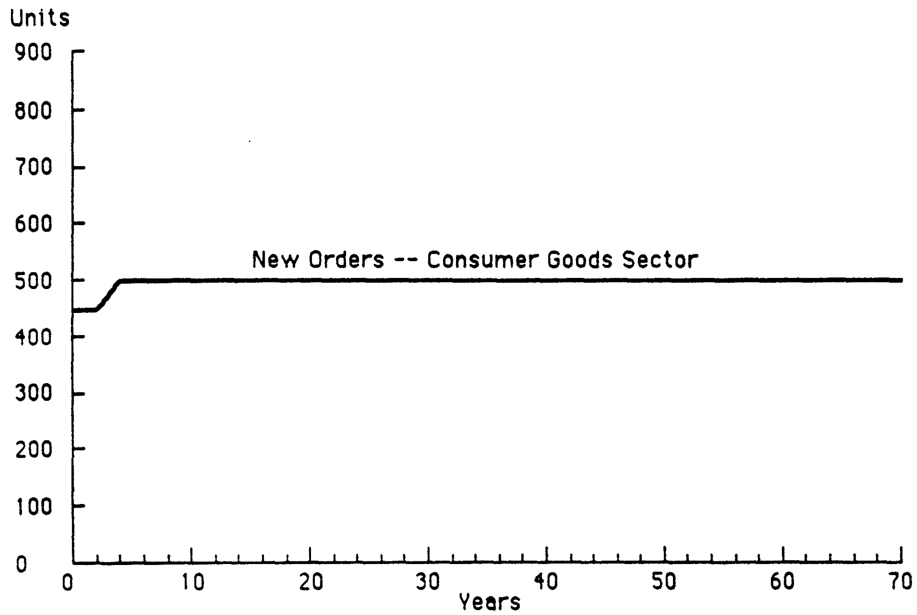

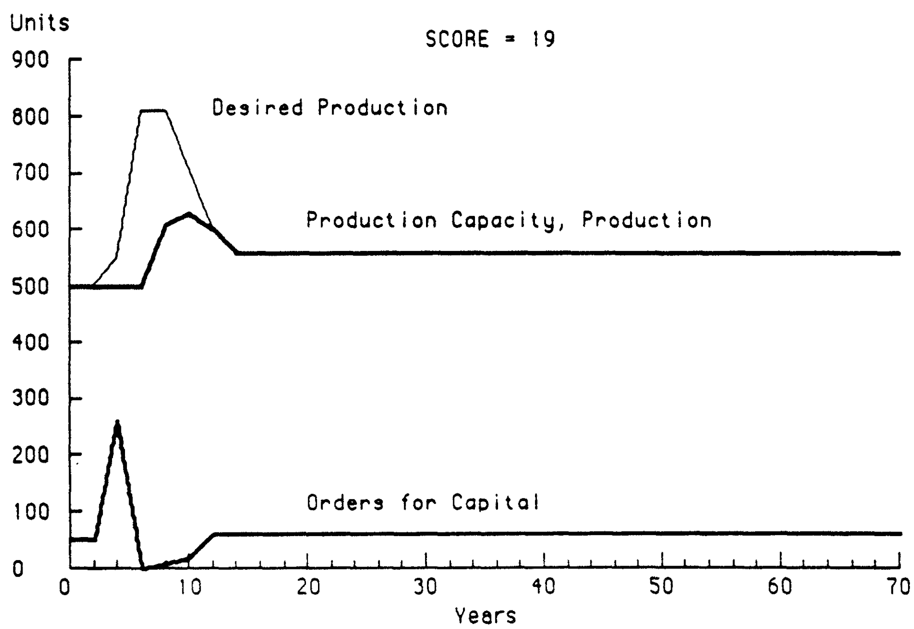

Figure 5. Typical experimental results. Note the large amplitude and long period of the cycles generated by the subjects. N.B.: vertical scales differ.

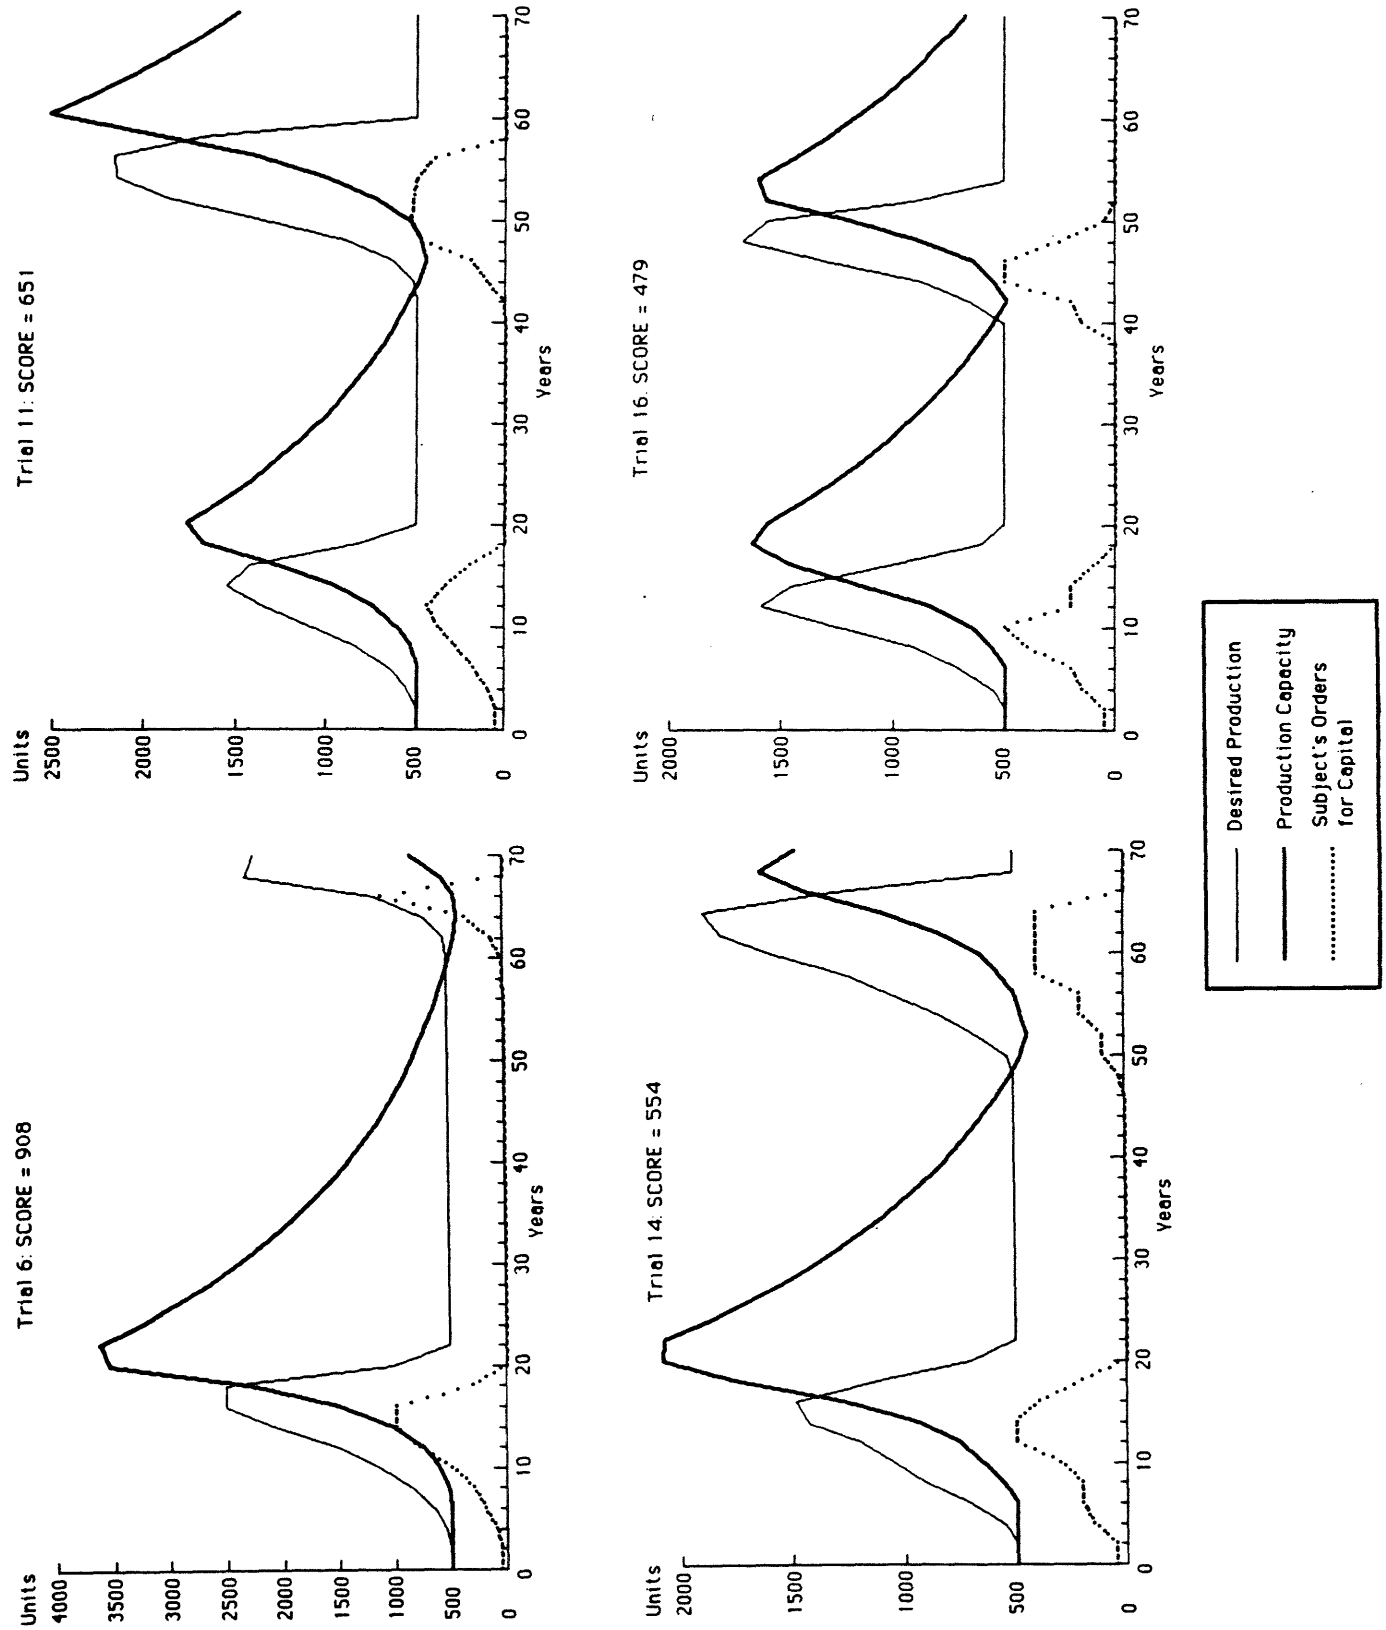

Figure 5 (cont.) Note: vertical scales differ.

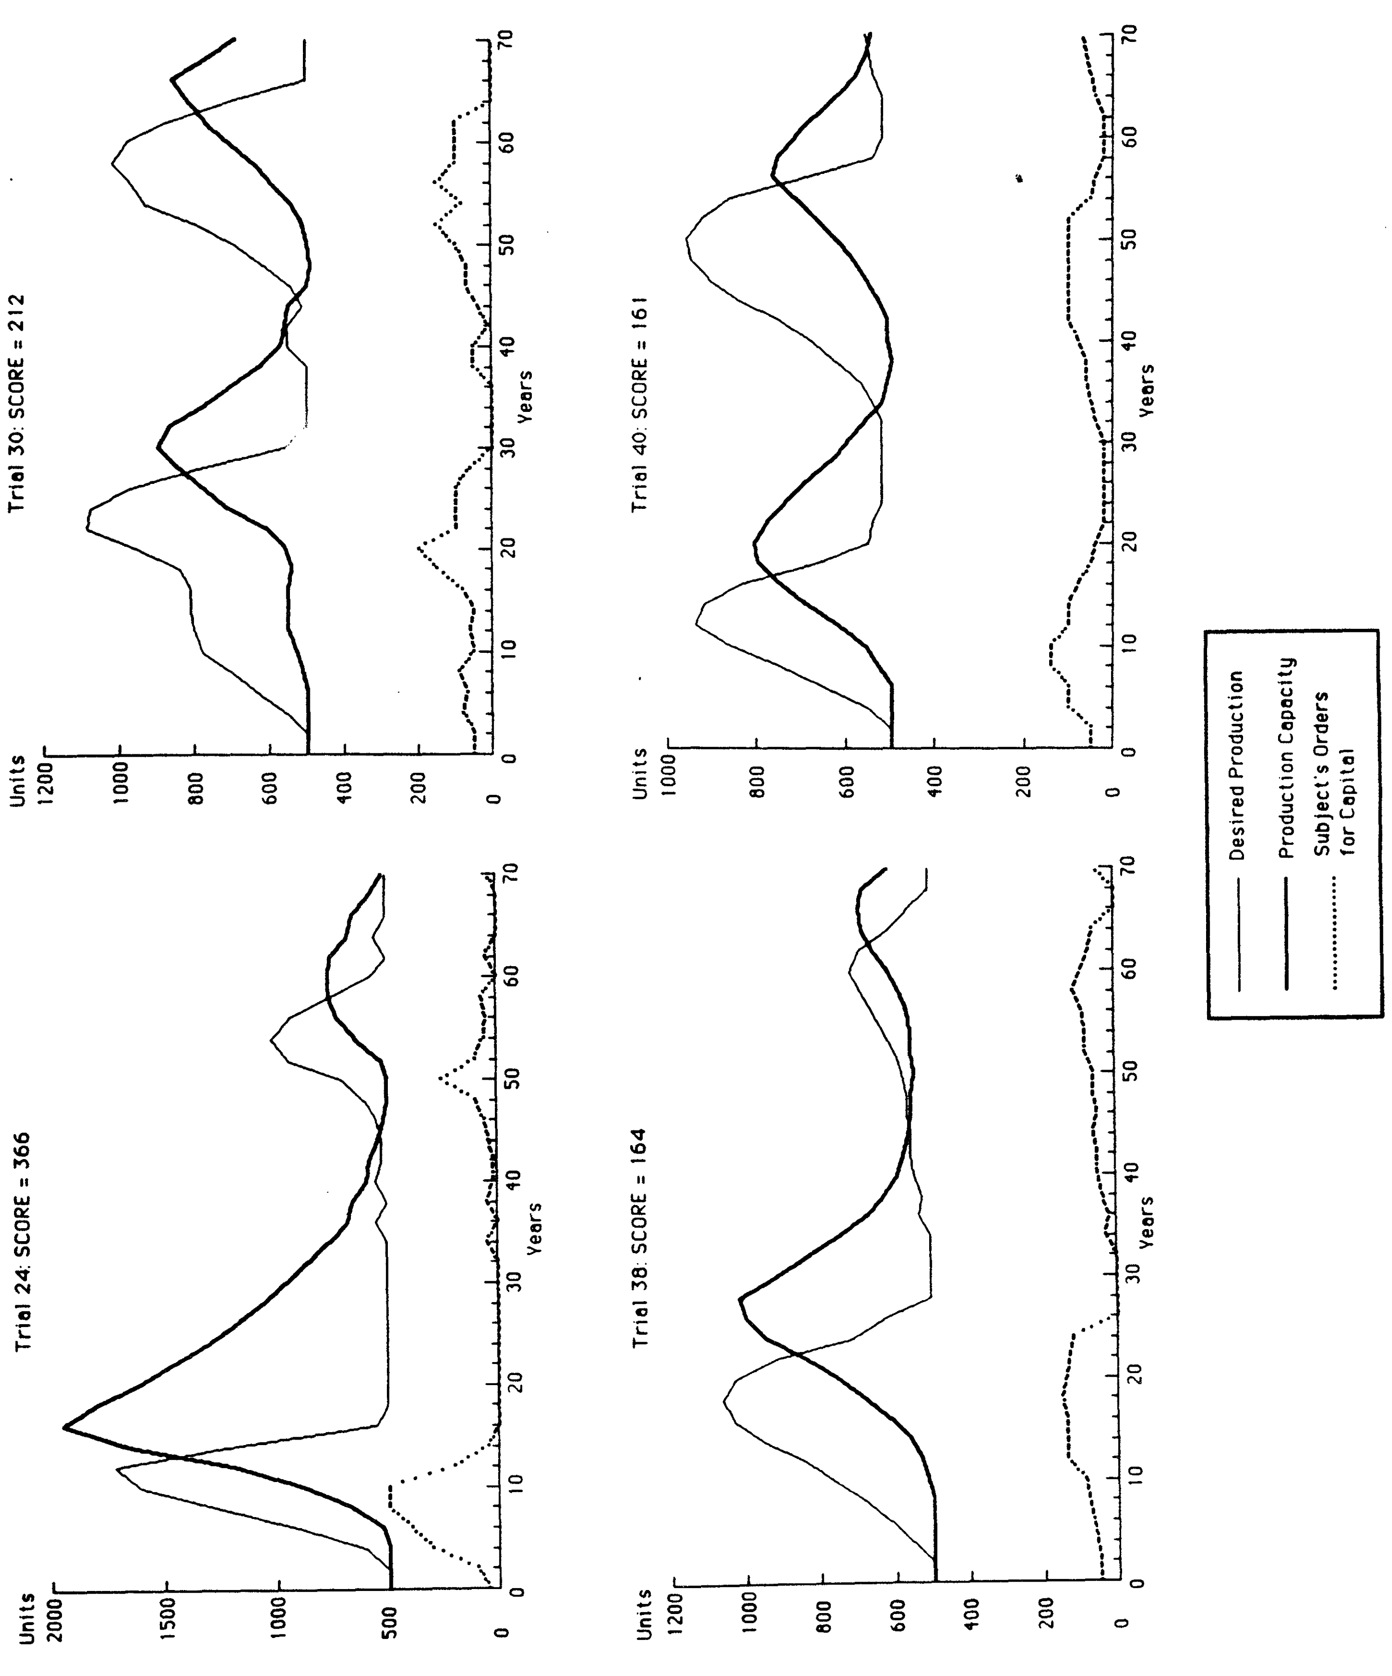

Table 1. Comparison of Experimental and Optimal Behavior

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

Numbers in parentheses exclude trial 1 as an outlier (score 8229; capacity peak >27,000; maximum order rate of 6000 units).

Table 2: Estimated parameters

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

Figure 6. Comparison of experimental and estimated orders for capital: Trials 5 and 38.

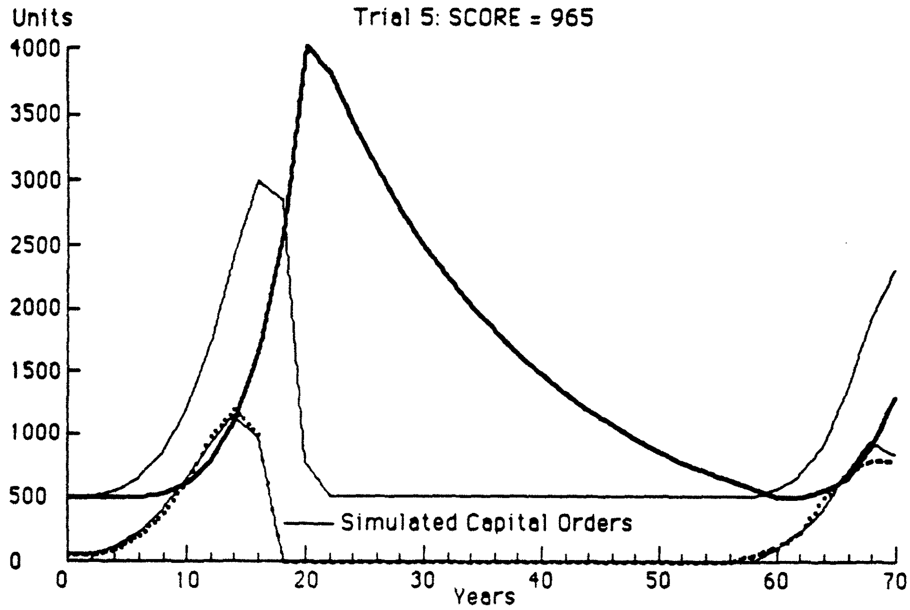

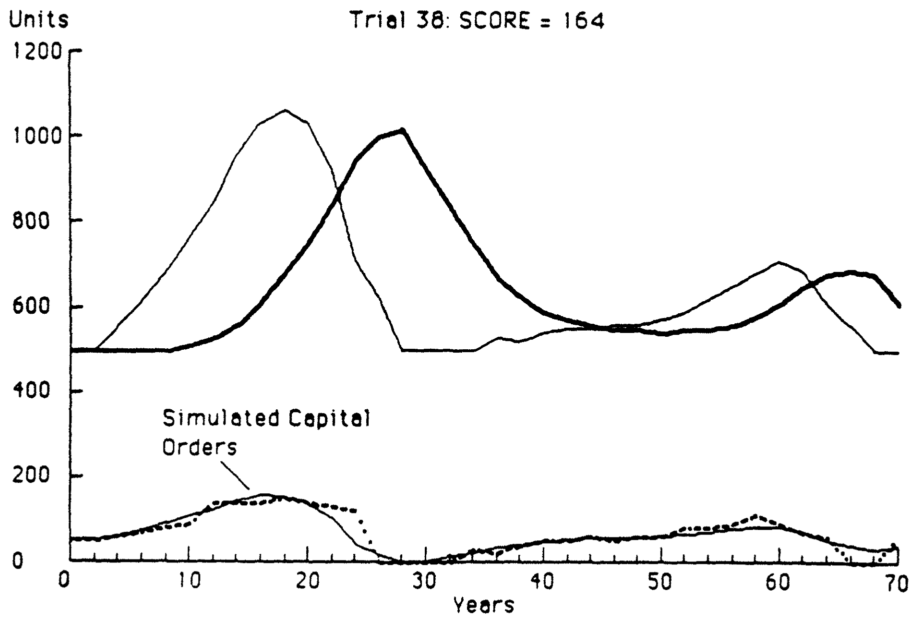

Figure 7. Comparison of experimental and estimated orders for capital: Trials 2 and 12.

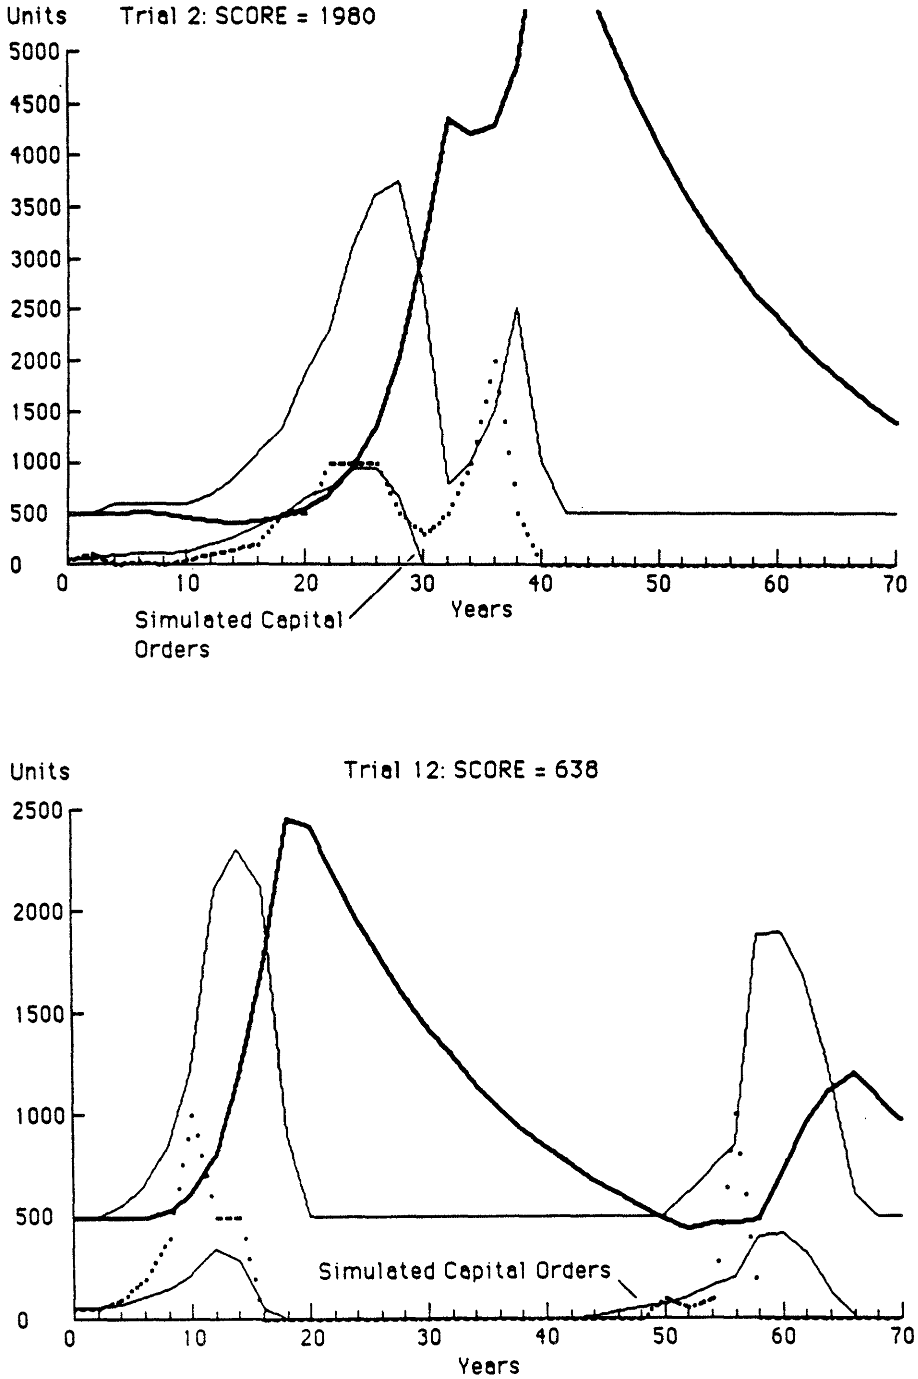

Figure 8. Testing the decision rule for intended rationality. Top: simulating the decision rule in the full model produces large-amplitude, long-period cycles similar to those produced by the subjects. Bottom: cutting the multiplier feedback means demand is exogenous and the capital acquisition delay is constant. The response is locally rational: there are no oscillations and equilibrium is restored rapidly. The same parameters ($\alpha_k = .55$ and $\alpha_{sl} = .40$) are used in both simulations.