Misperceptions of Feedback in Dynamic Decision Making*

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

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). 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 decision makers 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 themselves 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.

The system chosen for experimental investigation here is the multiplier-accelerator (MA) model of capital investment. First treated formally by Frisch (1933) and Samuelson (1939), MA models are central to many modern theories of business fluctuations (Goodwin 1951; Zarnowitz 1985 surveys recent theories). But while multiplier-accelerator (MA) models have been extensively studied, 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.

Traditional models such as those of Samuelson as well as their econometric descendents (see 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.

However, critical economists charge (correctly) that such decision rules are ad hoc, that they 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. In theory, all the variables and parameters affecting prices and quantities 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

---

* Condensed and reprinted by permission of the publisher from *Organizational Behavior and Human Decision Processes,* 43 (June 1989).

(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 (Simon 1979). The experiment described here offers the opportunity to test these theories of decision making directly.

METHOD

The experiment is based on a simple simulation model of the investment accelerator (Sterman 1985). 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.

The experiment is implemented on IBM PC-type microcomputers (disks for the PC or Macintosh are available from the author). A 'game board' is displayed on the screen (figure 1). 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 (representing two years) the subject decides how much capital to order. Details of procedure and the rules of the game are found in Sterman 1987. Subjects are responsible for one decision – how much capital to order and seek 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. The score indicates how well subjects balance demand and supply. Subjects are penalized equally for both excess demand and excess supply. Departures from the optimal score provide a simple metric for the "rationality" of the subjects' behavior.

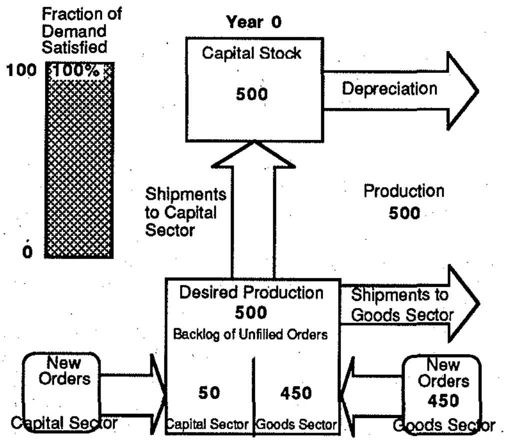

Figure 1. Computer screen as presented to the subject, showing the initial configuration. Color graphics and animation highlight the flows of orders, shipments, and depreciation.

The values of all system variables are displayed on the screen at all times. Subjects may examine 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.

The subject population (N=49) consisted of MIT undergraduate, master's and doctoral students in management and engineering, many with extensive exposure to economics and control theory; scientists and economists from various institutions in the US, Europe, and the Soviet Union; and business executives experienced in capital investment decisions including several company presidents and CEOs. All subjects were fluent in English. Sterman 1987 presents the experimental protocol and the equations of the model.

RESULTS

The trials were run for 36 periods. All were initialized in equilibrium with orders of 450 units/period from the goods sector and capital stock of 500 units. Capital discards are 10% per period, requiring the capital sector to order 50 units/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 the goods sector increases orders from 450 to 500, and they remain at 500 thereafter. The step input is not announced in advance.

The optimal response is shown in figure 2. Since the demand shock is unanticipated, capital sector orders remain at their initial level until *after* the demand shock. To reach the new equilibrium the order rate must exceed depreciation during the transient. Because capacity can only increase with a lag, the backlog of unfilled orders must rise above its 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. The optimal score is 19. Equilibrium is reestablished just 5 periods after the shock.

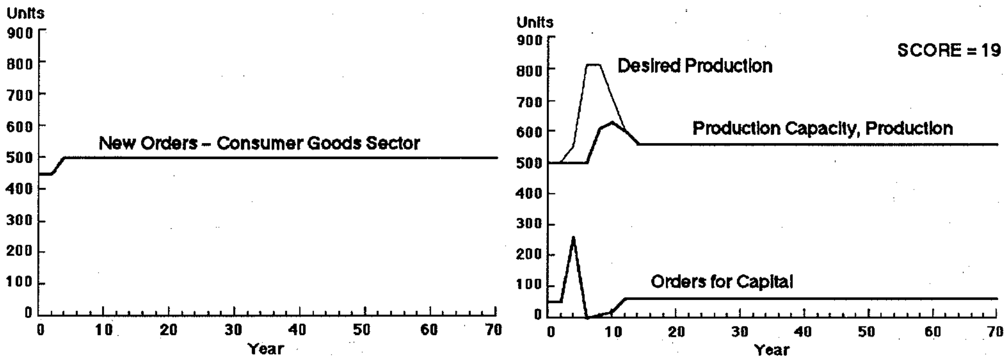

Figure 2. Each trial begins in equilibrium. In year 4 there is an unannounced increase in new orders placed by the consumer goods sector (left). The optimal response (right) returns the system to equilibrium by year 14.

The subjects behave quite differently. Figure 3 shows several representative trials; 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 via the multiplier, 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 capacity shortage slows the growth of capacity and frustrates 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 backlog is finally emptied. 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 more than 1600 units. Orders drop to zero. 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. The demand shock raises the total demand for capital by just 10%, but capacity rises over 300% at its peak.

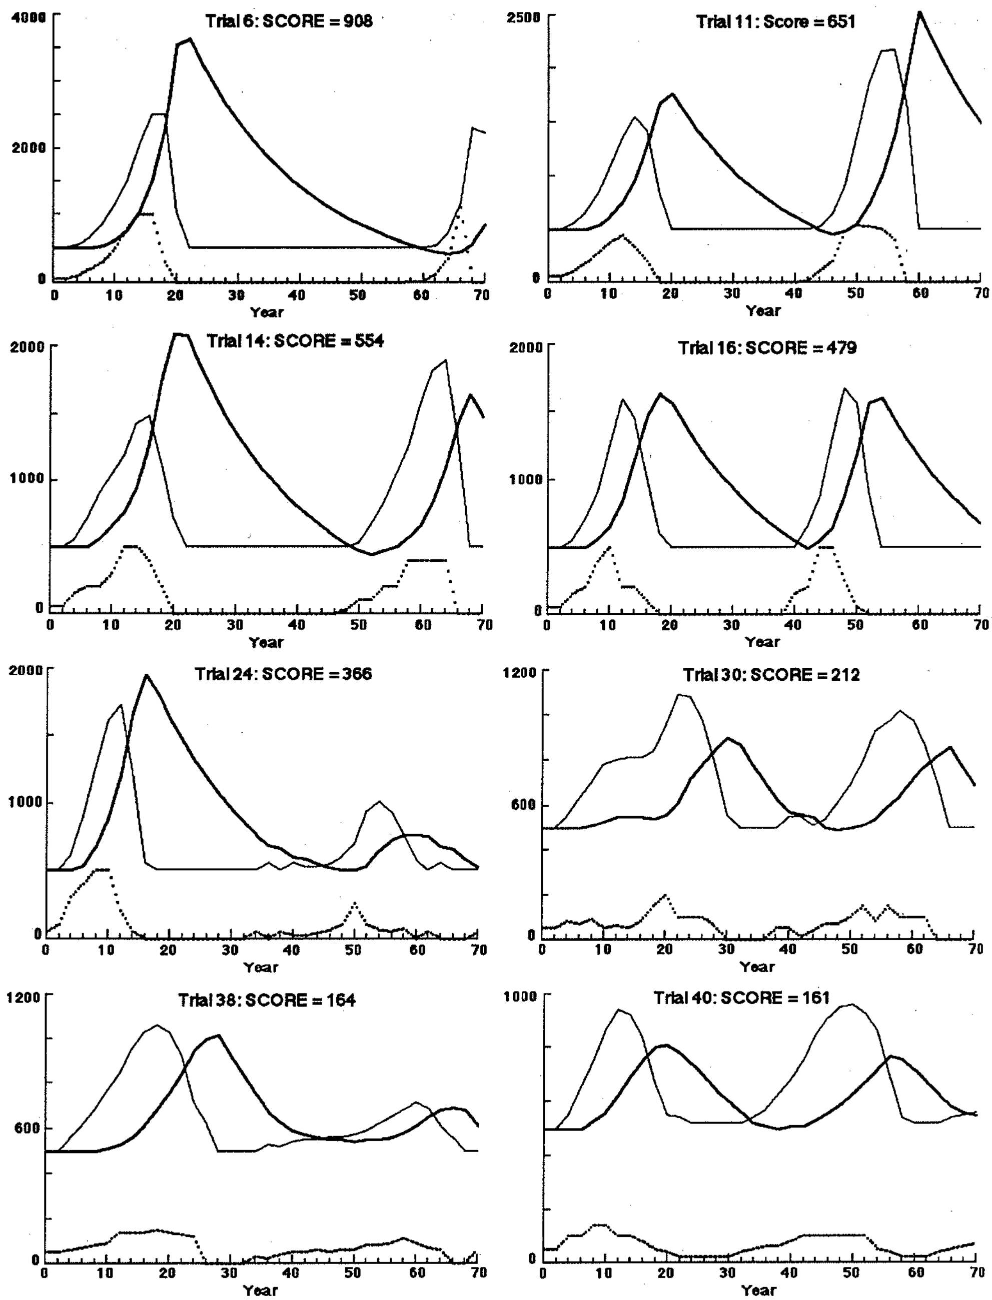

Figure 3. Typical experimental results. Note the large amplitude and long period of the cycles generated by the subjects. N.B.: vertical scales differ.

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.

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

	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

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 was used in the original simulation model upon which the experiment is based (Sterman 1985) and is a variant of rules long used in 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 DP, current production capacity PC, the rate of capital discards CD, the supply line SL of orders for capacity which the firm has placed but not yet received, and the capital acquisition delay CAD. 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:

$$CO_t = MAX(0, ICO_t). \tag{1}$$

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:

$$ICO_t = CD_t + AC_t + ASL_t . \tag{2}$$

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:

$$AC_t = \alpha_k \cdot (DK_t - K_t) \tag{3}$$

$$DK_t = \kappa \cdot DP_t. \tag{4}$$

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:

$$ASL_t = \alpha_{sl} \cdot (DSL_t - SL_t) \tag{5}$$

$$DSL_t = CD_t \cdot CAD_t \tag{6}$$

where DSL = the desired supply line and CAD is the capital acquisition delay. 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 (α_{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 receives the capital it requires to meet desired production.

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 easily shown to be the reciprocal of the fraction of demand satisfied 1/FDS (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:

$$CO_t = MAX(0, ICO_t + \varepsilon_t)\ ; \quad \varepsilon_t \sim N(0,\sigma^2) \tag{7}$$

and the parameters estimated by maximum likelihood methods, as described in Sterman (1989a).

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. The stock adjustment parameter α_k varies between .02 and 3.73 with a mean of .55; the supply line coefficient α_{sl} varies between 0 and 4.44 with a mean of .40.

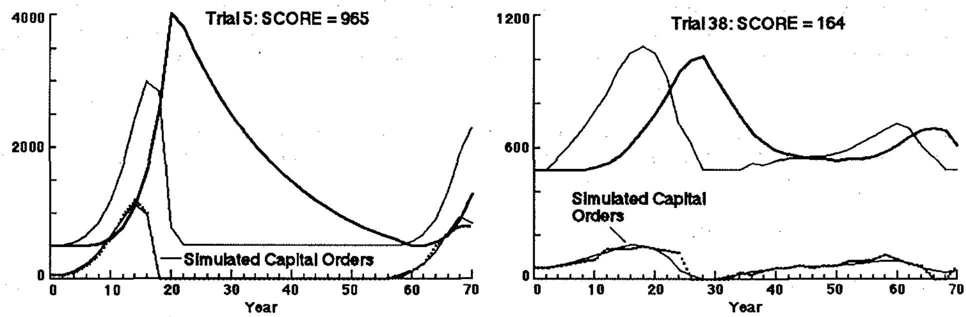

Figure 4. Comparison of experimental and estimated orders for capital: Trials 5 and 38.

To illustrate the performance of the rule, figure 4 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. Misunderstandings of the system structure or learning appear to account for the few trials for which the rule does not work well (Sterman 1989a).

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):

$$\text{Experimental Score}_i = \underset{(9.4)}{1.06} * \text{Simulated Score}(\alpha_k, \alpha_{sl})_i \qquad i = 2,...49; \qquad R^2 = .21 \qquad (8)$$

The slope of the relationship is highly significant and not statistically different from 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 decision maker seeks to maintain some stock or system state at a target level or within an acceptable range. The decision maker 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 many people apparently do not, causing instability).

There is no presumption that subjects calculated their decisions according to the equations of the rule. Yet clearly the rule is a good model of the heuristics they did use. Why did people behave in a fashion consistent with 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 2). 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 5 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 5a shows the full model as used in the experiment. The large overshoot of capacity, successive cycles, periodicity, and score are all characteristic of the experimental results. In figure 5b 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. 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 decision maker presumes it to be the response of the system to shocks is reasonable and appropriate.

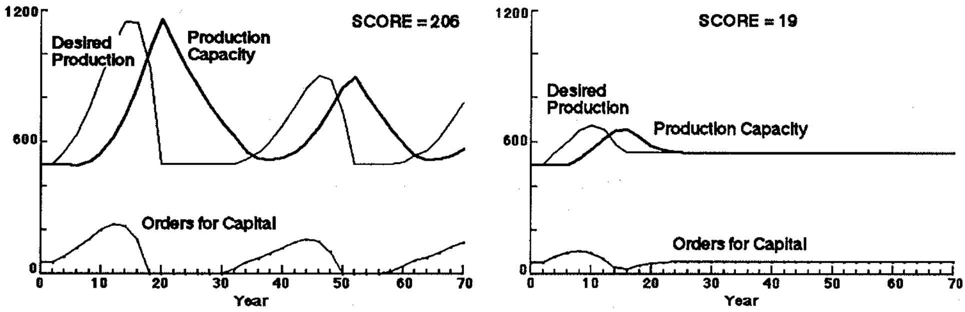

Figure 5a (left): simulating the decision rule in the full model produces cycles similar to those produced by the subjects. Figure 5b (right): cutting the multiplier feedback means demand is exogenous and the capital acquisition delay is constant. The response is locally rational: equilibrium is restored rapidly without oscillation. The same parameters (α_K= .55 and α_{Sl}= .40) are used in both simulations.

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 decision maker to assess correctly 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:

$$\ln(\text{Score}_i) = \underset{(42.8)}{5.3} + \underset{(5.1)}{1.7}*(\alpha_k)_i - \underset{(-3.7)}{1.1}*(\alpha_{sl})_i\ ; \quad i=2,...49; \quad R^2 = .43; \quad F = 16.8 \qquad (9)$$

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.

In fact, Sterman (1989b, 1988) shows that the estimated decision rule for approximately 20% of the subjects produce deterministic chaos when simulated. Consistent with the analysis above, the chaotic regime in parameter space exists in the region where capital stock adjustments are aggressive and supply line adjustments are weak.

2. *Misperception of feedbacks from decisions to the environment.* Figure 5 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 3). 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.

The interpretation above is supported by prior work in dynamic decision making, such as Doerner (1980) and Kluwe et al. (1984). Though these experiments employed rather different tasks, both concluded that subjects tend to think in single-strand causal *series* and thus have difficulty in systems characterized by causal *nets* (i.e. side effects). Broadbent and Aston (1978) likewise found that managing an econometric model of the U.K. economy produced little change in subjects' (verbally reported) understanding of economic relationships. With sufficient experience, however, subjects were able to control the simulated economy better than initially. The results here reinforce these findings and suggest that performance is degraded still further in systems characterized by causal *loops*, time delays, and nonlinearity, a result consistent with Brehmer's (1986) analysis of a fire-fighting simulation.

Do such misperceptions of feedback exist in the real world, or are they artifacts of the unfamiliar task of the experiment? 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, and the time required for learning may exceed the tenure of individual decision makers. 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.

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 decision-making 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."

CONCLUSIONS

The results of this work have several implications for research in dynamic decision making and economics. Traditional macroeconomic models of investment behavior assume individual firms follow a difference-reduction heuristic. Modern theories assume firms behave so that their behavior is optimal with respect to some intertemporal objective function. The experimental results show 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 decision making exacerbates instability.

Finally, it appears that the experimental exploration of dynamic decision-making strategies in aggregate systems is feasible. The fidelity and flexibility of simulation models enables the investigator to construct rich, complex decision-making 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 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.

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 Decision Making," presented at the 1986 annual meeting of the Judgment/Decision Making Society, New Orleans, November.

Broadbent, D. and B. Aston (1978). "Human Control of a Simulated Economic System," Ergonomics. 21, 1035-1043.

Doerner, D. (1980). "On the Difficulties People Have in Dealing with Complexity," Simulation and Games. 11, 87-106.

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.
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 Decision Making 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.
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.
Kluwe, R., C. Misiak, and H. Reimann (1984). "Lernvorgange beim Umgang mit Systemen: Die Ausbildung subjektiver Ordnungsstrukturen durch Erfahrung mit umfangreichen Systemen," Technical Report 7, Fachbereich Padagogik, Abteilung Allgemiine Psychologie, Hochshule der Bundewehr, Hamburg.
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.
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.
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.
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.
Samuelson, P. (1939). "Interactions Between the Multiplier Analysis and the Principle of Acceleration," The Review of Economic Statistics. 21, 75-78.
Simon, H.A. (1979). "Rational Decision Making 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.
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. 33(12), 1572-1592.
Sterman, J.D. (1988). "Deterministic Chaos in Models of Human Behavior: Methodological Issues and Experimental Results," System Dynamics Review. 4, 148-178.
Sterman, J.D. (1989a). "Misperceptions of Feedback in Dynamic Decision Making," Organizational Behavior and Human Decision Processes. 43.
Sterman, J.D. (1989b). "Deterministic Chaos in an Experimental Economic System," Journal of Economic Behavior and Organization. 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.
Zarnowitz, V. (1985). "Recent Work on Business Cycles in Historical Perspective: A Review of Theories and Evidence," Journal of Economic Literature. 23, 523-80.