Effects of Feedback Complexity on Dynamic Decision Making

John D. Sterman
and
Ernst Diehl

WP# 3608-93-MSA August, 1993

D-4401

Effects of Feedback Complexity on Dynamic Decision Making

Ernst Diehl and John D. Sterman

Sloan School of Management
Massachusetts Institute of Technology

August 1993

We thank Bent Bakken, John Carroll, Christian Kampmann, Don Kleinmuntz, Paul Andreassen, and seminar participants at MIT for assistance and comment. This work was supported in part by the Organizational Learning Center, MIT Sloan School.

Please direct correspondence to John Sterman at Sloan School of Management, Massachusetts Institute of Technology, 50 Memorial Drive, E53-339, Cambridge, MA 02142 or jsterman@mit.edu.

Effects of Feedback Complexity on Dynamic Decision Making

ABSTRACT

Prior research shows people suffer from misperceptions of feedback in dynamic settings, generating systematic dysfunctional behavior in the presence of dynamic complexity – settings with multiple feedback loops, time delays, and nonlinearities. However, prior work has not adequately mapped the effect of these elements of complexity on performance. We report an experiment where subjects managed an inventory in the face of stochastic sales, a classic dynamic decision task. We vary the time delays and strength of the feedback loops as treatments to explore the impact of these elements of dynamic complexity on behavior. Subjects faced financial incentives and had extensive opportunities to learn. Yet performance was significantly worse than optimal across all conditions. Subjects outperformed a naive 'do-nothing' rule in the simple conditions, but performance deteriorated dramatically with increasing time delays and feedback effects, and most were outperformed by the do-nothing rule in the complex conditions. Regression analysis of subjects' decisions showed most ignored the supply line of pending production and undercontrolled the system. Undercontrol increased significantly with growing time delays and feedback strength, showing subjects were insufficiently adaptive despite perfect knowledge of system structure and parameters. Subjects' understanding of complex feedback settings declines as delays between cause and effect increase, and as actions have stronger side effects. Few indications were found of active experimentation or learning: the need to control seemed to override the ability to learn.

KEY WORDS: Decision making, simulation, feedback, experimental economics, system dynamics

Dynamic decision tasks arise whenever decisions made today alter the state of the system and thus the information that conditions decisions to be made tomorrow. Dynamic decision tasks are common, indeed inescapable, in many domains. Driving a car, managing a firm, controlling a chemical plant and controlling the money supply are all dynamic tasks because the decision maker and the system are entwined in feedback loops whereby decisions alter the environment, giving rise to new information and leading to new decisions (Forrester 1961, Richardson 1991). Dynamic decision tasks vary in terms of the *dynamic complexity* of the system. Many real tasks are dynamically complex because they involve time delays, stocks and flows, multiple feedback processes and nonlinearities. Prior research shows people have great difficulty managing dynamically complex tasks, generating significant, systematic, and costly errors (Sterman 1989a, 1989b; Paich and Sterman 1993).

Sterman (1989a, 1989b) argued that the observed dysfunction in dynamically complex settings arises from systematic 'misperceptions of feedback' – that the mental models people use to guide their decisions are dynamically deficient. People generally adopt an event-based, 'open-loop' view of causality, ignore feedback processes, fail to appreciate time delays between action and response and in the reporting of information, do not understand stocks and flows, and are insensitive to nonlinearities which may alter the strengths of different feedback loops as a system evolves. In addition to the studies cited above, the misperceptions of feedback (MOF) hypothesis is generally supported by studies in experimental economics, psychology, and management (Smith, Suchanek and Williams 1988, Funke 1991, Brehmer 1992).

For example, Brehmer (1990) developed a computer simulation of a forest fire in which subjects played the role of fire-fighting chiefs, deploying field units from a headquarters in the rear. The task included significant feedback complexity, particularly the self-reinforcing (or *positive*) feedback by which the fire spreads. Brehmer found that subjects' ability to control the simulated fire was poor, and the rate of learning slowed significantly, when there were time delays between action and response (the deployment of fire fighting units in the field) and in the reporting of information (receipt of reports on the location and status of the field units).

Similarly, Sterman (1989b) examined a simple inventory management task, the 'beer distribution game,' in which subjects sought to minimize costs as they managed the production and distribution of a commodity. Though simplified compared to real firms, the task was dynamically complex as it included multiple feedbacks, time delays, nonlinearities, and accumulations. Average costs were ten times greater than optimal. The subjects generated costly oscillations with consistent amplitude and phase relations, even though the demand for the product was essentially constant. Econometric analysis of subjects' decisions showed that people were quite insensitive to the time delays in the system. In particular, people did not account well, and often not at all, for the supply line of orders that had been placed but not yet received, causing them to order many times more beer than they actually needed. Sterman (1989a) found subjects exhibited the same phenomena in a simulated macroeconomy with time delays and feedback loops; here subject performance was 19 times worse than optimal. In both experiments the deviations from optimality were systematic – subjects generated large, persistent oscillations with characteristic amplitude and phase relations among the variables. In both experiments simulation of the decision rules estimated for the subjects showed that approximately one third were intrinsically unstable, so that the system never reached equilibrium. About one-quarter of the estimated rules yield deterministic chaos (Sterman 1989c, 1988). The heuristics people used interacted with the feedback structure of these systems to yield severe, systematic, persistent, and costly oscillations.

However, recent studies are limited in several respects. Many report the results of first trials where subjects had little opportunity to learn (Hogarth 1981). In others, critical information is not available to the subjects. For example, in Brehmer's fire fighting task subjects were not informed of the delay in the reporting of the status of field units, though they could infer it from the information on the screen. In the beer distribution game the supply line of orders is not available to the subjects but must be inferred from the history of orders and deliveries. An additional problem arises from the sheer complexity of these tasks. Though the experimental tasks reported in the literature are simplified compared to reality, it is hard to disentangle the roles of the different types of feedback complexity. In some studies the elements of feedback complexity were not varied as treatments, while in others, only a subset were varied, or several were varied simultaneously. While the rich variety of dynamic decision making tasks has revealed much about the flaws in people's abilities to manage dynamic complexity, this same richness has hindered comparability of results and the ability to reach general conclusions about the relationship of task structure to dynamic decision making performance, a point made as well by MacKinnon and Wearing 1985.

In this study we map the effects of time delays and feedback processes on decision making and performance in a dynamic task. By separately varying time delays and feedback effects we can disentangle the roles of these two elements of dynamic complexity. Clearly there are many dimensions of complexity beyond time delay and unintended feedback. However, time delays and feedback processes are fundamental and pervasive in dynamic decision environments. Understanding their role lays a foundation for further studies of the interactions between dynamic complexity and human performance (see also Kampmann and Sterman 1992, Bakken 1993).

DYNAMIC DECISION ENVIRONMENTS

A vast number of human activities can be characterized as attempts to control a stock and maintain its value close to a target value or within an acceptable range (Sterman, 1987, 1989b). Stock adjustment is the prototypical dynamic decision making task. Stocks (state variables) are accumulations of their various inflows and outflows, and thus represent the memory in a dynamic system by which past events condition the state of the system that leads to new decisions. Stock adjustment problems are prevalent on different levels of aggregation. People change their car's velocity to drive at a desired speed, regulate the water's temperature to shower comfortably, and vary diet and exercise as they seek to maintain a desired weight. Firms set production schedules to control inventories, hire and fire employees to meet their labor needs, and borrow money to manage their cash balances. A chemical plant operator adjusts flows of inputs to keep a process operating in the safe range. The federal reserve engages in open market operations to adjust the stock of money in the economy towards target ranges. In all of these examples, the objective is to maintain a stock at its target value in the presence of disturbances such as losses, usage, or decay.

In general, a decision maker's actions do not directly affect the flows into or out of a stock.

Rather, the decision maker typically controls flows only after delays. Firms cannot acquire new plant and equipment instantly, but must order new equipment from suppliers or build new facilities. While awaiting delivery the orders accumulate in a supply line of plant under construction or equipment on order. There may be multiple influences, both exogenous and endogenous, on the flows that affect the stock to be managed. The breakdown and discard of existing equipment – the outflow from the stock of equipment – is partly exogenous and stochastic, and partly an endogenous function of how intensively the equipment is used and how well it is maintained. Similarly, the rate at which new equipment is delivered depends partly on exogenous factors such as the time required to transport the equipment from the supplier to the customer and partly on the actions of the customer itself. For example, if a customer places enough orders, they may exceed the supplier's capacity, slowing delivery as the orders wait in a backlog until they can be processed.

The decision makers' task is to set an appropriate order rate so as to keep the stock close to the desired or target value, taking into account the likely outflows from the stock, the time delay between placing and receiving an order, and the various feedbacks that might alter conditions in response to the decision maker's own actions (figure 1). The decision maker must attend to the value of the stock compared to its target, creating a self-correcting, or negative, feedback loop. Because of the time delay between the initiation and completion of these control actions, this negative loop is potentially oscillatory.[1] To compensate a wise decision maker may attend to the supply line of pending production, creating a second negative loop to prevent overordering (though the evidence shows many do not). Finally, the decision maker may try to anticipate the likely outflows from the stock (a feedforward or forecasting component). To the extent decision makers can forecast the future flows affecting the stock, they can order in advance and prevent the stock from deviating from its target value.

The experimental system: In the present experiment we created a generic stock adjustment task with the basic features identified above. The task is extremely general in structure, though we gave it a business context for concreteness (Diehl 1992 provides complete documentation). Subjects manage the production of a firm (figure 1) and seek to minimize their cumulative costs throughout each trial. Costs arise from discrepancies between inventory and its target value and from the adjustment costs of changing production. Subjects can review each period their inventory levels, sales, production, and costs. They decide how much to change production and enter their decision. The revised rate of production adds to the work in process inventory (the contents of the production delays, if any) until the goods are completed and added to inventory. Sales decrement inventory. Time advances to the next round, and they make their next production decision.

More formally, the task is described by the following equations. The stock to be managed is Inventory, I. Inventory is increased by production, P, and decreased by sales, S:

$$I_t = I_{t-1} + P_t - S_t. \quad (1)$$

Production is determined by production starts, P*, after a delay of δ periods to represent the time required for the manufacturing process:

$$P_t = P^*_{t-\delta}. \quad (2)$$

In the experiment subjects determine the change in production starts, Δ:

$$P^*_t = P^*_{t-1} + \Delta_t \quad (3)$$

$$\Delta_t = \text{<Subject's decision>} \quad (4)$$

When there is a delay, production that has been initiated but not yet completed accumulates in a supply line of work-in-process inventory, given by ΣP^*_{t-i}, $i = 1,\ldots,\delta$. The length of the production delay was a treatment variable in the experiment.

The outflow of inventory is sales (on the right side of Figure 1). The demand formulation was designed to test the ability of people to understand and manage systems with multiple feedback processes. In dynamically complex systems a control action may have unintended side effects. Side effects may reinforce or oppose the intended effects of the decision, and may operate with different delays. In the context of the manufacturing setting used here, side effects arise from the linkages among firms in the broader community. For example, the more goods produced in an economy, and thus the higher the employment, the greater are people's incomes. Consumers use their higher income to purchase more goods and services, boosting the sales of businesses throughout the community. Thus as a firm increases production, it indirectly increase its own sales. Higher sales further deplete inventories, leading to pressure for still higher production, in a positive (self-reinforcing) feedback known in macroeconomics as the 'Keynesian multiplier'. Side effects may also form negative feedback loops. In the business context used here such an effect would mean that rising production leads to falling sales, and vice-versa. For example, suppose by increasing your production you signaled opportunities for expansion to your competitors, who respond with price cuts or marketing programs that lower your demand. Here rising production leads to declining sales, unintended inventory accumulation, and pressure to cut back production, forming a negative or self-correcting feedback loop. To capture such side effects, sales in our model consist of an exogenous component, X, and a dependent or endogenous component, N:

$$S_t = X_t + N_t \tag{5}$$

The endogenous component of demand, N, is determined by the current production rate and the gain, γ, of the side-effect feedback:

$$N_t = \gamma P_t \tag{6}$$

Thus, for example, a gain of +.3 indicates that a 10% increase in production causes an additional 3% increase in the endogenous component of sales. The gain (sign and strength) of the feedback from production to sales was the second treatment variable in the experiment.

Finally, the exogenous component of sales follows a random walk :

$$X_t = X_{t-1} + U_t \tag{7}$$

where the changes, U, are drawn from a uniform distribution with mean zero and range from -15 to +15 units/period. The random walk provides a challenging input to the system, one that requires subjects to control the system actively or face large, persistent, and costly deviations of their inventories from the target value.

Costs are determined by both the inventory position of the subjects and the rate of change in production. Inventory costs capture the cost of deviations of the stock from its target or set point. The set point for inventory was zero to simplify the calculations required of the subjects. In general, stock adjustment tasks involve adjustment costs or costs of control effort as well as costs of deviations between the stock and its desired value. In our context, adjustment costs arise from changes in production and can be thought of as the costs of hiring and firing workers, expediting materials acquisition, or subcontracting the work. A quadratic cost function is assumed. Quadratic costs are reasonable approximations to the loss functions in many stock management settings, including inventory management (Holt et al. 1960) and allow analytic solutions to the optimal production problem to be computed conveniently. Thus costs, C, are given by

$$C_t = aI_t^2 + b\Delta_t^2;\ a=1,\ b=2. \tag{8}$$

If the adjustment costs are too small, the optimal solution is to jump immediately to the optimal inventory level, leading to unrealistically large changes in production. If the adjustment costs are too large, the optimal solution is unrealistically slow correction of inventory discrepancies. Simulation studies (Diehl 1989) showed that the coefficients a=1, b=2 represent a good tradeoff between inventory and adjustment costs.

HYPOTHESES

The central issues are (1) the extent to which subjects are able to control the system in the face of dynamic complexity; and (2) the extent to which their decision making behavior is sensitive to different elements of complexity. That is, how does performance vary as the feedback complexity of the task increases? What are the differential effects of delays and feedback effects? How do delays and feedback processes interact to influence performance? Are subjects sensitive to the presence of delays and feedbacks? Do they alter their use of information and decision making heuristics appropriately as the feedback structure of the task changes?

The treatment variables in the experiment are the length of the production delay, δ, and the sign and magnitude of the gain of the side effect feedback, γ. The values of delay and feedback gain chosen as treatment levels spanned a wide range. We selected delay lengths $\delta = 0, 2,$ and 4 periods and gains of $\gamma = -0.6, -0.3, 0, +0.3,$ and $+0.6$, yielding a 3 (delay) by 5 (gain) factorial design. The zero delay, zero gain condition corresponds to a simple task such as considered by MacKinnon and Wearing (1985). Here there are no delays between your decision and its realization, nor are there any side effects. Essentially, your task is to keep your bathtub filled to the right level by adjusting the tap, a straightforward task in which performance should be excellent. When there is a delay, even the optimal rule bears additional costs since an inventory discrepancy cannot be corrected until δ periods have passed. To the extent subjects fail to account for the delay, they are likely to boost production throughout the period during which inventory remains too low, leading to excess inventory, oscillation, and costs higher than optimal.

The anticipated role of the side effect feedback is more subtle. Suppose the gain of the side effect loop is positive. Every increase in production then causes a proportional increase in sales. If inventory is too low and the decision maker increases production, sales will rise, making it harder to close the gap, and leading to more pressure to produce. Conversely, cuts in production to reduce excess inventory cause sales to fall, slowing the reduction of inventory below the intended rate. The side effect forms a positive, self-reinforcing feedback that undercuts the stabilizing effect of the main control loop. To achieve a given change in inventory it is necessary to make larger and more costly changes in production. Thus the positive side effect raises optimal costs. To the extent managers fail to account for the effect of the positive side effect loop, they will undercontrol the system, leading to larger and more persistent inventory discrepancies.

In contrast, negative gain is highly stabilizing and reduces optimal costs. When the side effect gain is negative, every increase in production produces a corresponding decrease in sales. If inventory is too low and the decision maker increases production, sales will fall, helping to raise inventory and closing the gap more quickly and more cheaply. Similarly, production cutbacks to reduce excess inventory induce a sales increase, speeding the adjustment of inventory to its target level. The side effect forms a negative feedback that assists the stabilizing effect of the control loop formed by attempts to regulate inventory. To achieve a given change in inventory requires smaller and less costly changes in production. Thus the negative side effect loop is highly stabilizing and lowers optimal costs. However, to the extent subjects fail to account for the stabilizing effect of the negative gain they will overcontrol the system, leading to frequent overshoot and oscillation.

When both production delays and side effects are present the control task becomes considerably more difficult. Long delays and high positive gain are highly destabilizing, as the delay makes it hard to correct inventory gaps and the positive feedback undercuts the corrective effect of production changes. The production delay means that the impact of the side effect is no longer contemporaneous with the impact of production, increasing the likelihood of inventory overshoot and oscillation if subjects are insufficiently sensitive to either the delay or feedback. Likewise, under long delay and negative gain, the stabilizing impact of the side effect loop is out of phase with the production decision. Failure to account for the delay or the negative loop will lead to instability and higher costs.

At one extreme, perfectly rational subjects would always use the optimal decision rule. The weights they would apply to the available information would change appropriately as the delay and gain conditions changed. Performance would diverge from optimal only to the extent subjects applied the optimal rule inconsistently. At the other extreme, a subject might choose to take no action, a 'no-control' rule. A no-control strategy minimizes adjustment costs and cognitive effort but allows inventory to follow a costly random walk. We use these two extreme strategies as benchmarks for assessment of subject performance.

An extreme interpretation of the misperceptions of feedback (MOF) hypothesis suggests that subjects would be completely unresponsive to the presence of time delays and feedback loops. Subjects under this extreme MOF hypothesis would make decisions as if there were no delay and no side effect loop regardless of the actual delay and gain condition. Estimates of subjects' cue weights would reveal heavy reliance on inventory but no weight on the supply line (future values of production). As a result, performance would be significantly worse, *relative to optimal,* in conditions with long delays and high positive gain, and better in conditions with no delay and no or negative gain.

However, the task is extremely simple. There is no overt time pressure. The delay and gain conditions are highly salient: the concepts of delay and gain are emphasized in the task briefing; the delay and gain values appear at all times on the information display; and the values of the supply line (future production for the next δ periods) and the impact of the side effect feedback (dependent sales) are clearly displayed at all times. It would be remarkable if subjects were completely unresponsive to these changes in task structure. Indeed, even in Sterman's (1989b) beer game experiment, where the supply line was not presented to the subjects but had to be constructed from other cues, a minority of subjects were able to account for the supply line. However, even if subjects are aware of the task structure, the MOF hypothesis suggests they may not fully understand the implications of delays and feedbacks. In particular, their ability to infer the future dynamics of the system from their causal map or mental model may be inadequate in the presence of delays and feedback loops. Thus a less restrictive interpretation of the MOF hypothesis suggests that subjects will adjust their decision making heuristic somewhat as the delay and feedback conditions change. To the extent subjects' mental models underemphasize the importance of delays and feedback, and to the extent subjects have difficulty inferring the consequences of these elements, such adjustments are likely to be insufficient, especially in conditions with long delay and strong positive feedback. As a result, performance should be relatively worse compared to optimal in the difficult conditions (long delays and high positive gain) and comparatively better in the easy conditions (no delay, no or negative gain).

METHOD

Design: The three delay and five gain levels define a factorial design with 15 conditions. Since performance in the task is likely to improve with experience, each subject should play many times. A Latin square was used to assign the 15 experimental conditions to the subjects and trials such that every subject received every treatment, and every treatment appeared in every position in the sequence (figure 2), yielding 225 cells.

Fifteen different realizations of the random walk in exogenous sales (eq. 7) were used. If the same pattern were reused, subjects would gradually develop a good forecast of demand, obviating the need to control inventories and confounding changes in their stock management heuristic over trials with changes in the quality of their demand forecasts. Using different patterns ensures subjects face unintended changes in inventory, requiring them to take control actions and revealing their stock management heuristics. On the other hand, using different random walks in all 225 cells would add considerable variance to the results. To minimize the error variance between conditions every subject received the same sequence of random-walk patterns. As a result the sequence of demand patterns is confounded with any learning or practice effects. However, a strong practice effect is expected in repeated decision tasks of this type, and the focus of the study is the subject's heuristics and their dependence on the treatments, not the practice effect. Thus the inability to separate the practice effect from the sequence effect does not compromise the results.

Participants: Seventeen MIT students, fourteen undergraduates and three graduate students, enrolled to participate in the study. Two subjects, for undetermined reasons, did not complete the study: One terminated after the first session and the other terminated after the second session. Their results indicated no large differences from other subjects.

Procedure: Subjects were paid for their participation in the study. Subjects' total pay consisted of a base pay of $20 plus a performance-based amount. Subjects were informed that for the computation of their reward only their 12 best games would count. However, all data were utilized in the statistical analysis. To minimize fatigue effects the trials were run in four sessions, with three trials in the first session and four in the sessions to follow. Participation required, on average, six hours spread over four sessions in a two to three week period. Actual performance-based pay varied between $5 and $45. On average, subjects received a total of $40.

Subjects were provided with an extensive written description of the task covering the structure and context, including how all variables were related and calculated (Diehl 1992). Subjects were fully informed about the nature of the disturbances to the system. In particular, the endogenous and exogenous components of demand were explained, and they were informed that the exogenous component was a random walk. The concept of a random walk was explained, and subjects were told how to forecast it optimally. Specifically, they were told "your best bet is to expect that independent sales in the next period will be the same as they are in this period."

The task was implemented on Macintosh computers using a spreadsheet display (figure 3). The subjects, MIT students, had high computer skills and were familiar with such displays. The screen provided complete information on the current state of the system, including production, the dependent and independent component of sales, total sales, the change in inventory, the inventory level, and current and accumulated costs. In conditions with non-zero production delay the display also showed the supply line of work in progress which determines production for the next δ periods. To assist subjects, the current delay and gain conditions were always displayed. The computer recorded the time between decisions as well as the decisions themselves.

Two emulated practice rounds were given in order to familiarize subjects with the information display and task. The practice rounds consisted of a paper copy of actual computer screens and two hypothetical examples. Subjects were informed that in the experiment the computer would calculate the consequences of their decision, but that for practice purposes the subject would perform the computations. Then the subjects were encouraged to make a decision for production and compute the consequences of the decision by hand. The experimenter provided them with the value of independent sales. Subjects then proceeded to compute the rest of the variables. Erroneous computations were immediately corrected and an explanation provided. Most subjects did not make any errors in the second practice task.

After the practice task, the instructor repeated the rules of the experiment. Subjects were informed that they would play 15 games altogether, each of 32 periods. They were told that the production delay and strength of feedback to dependent sales might vary from game to game. They were further informed that some conditions were intrinsically more difficult than others and were told to not be discouraged by what they might consider extremely high costs. Specifics about what high costs might be were not provided.

The computer was then turned on and the subject was asked if he or she had ever used a

Macintosh before; if not, these mechanics were explained. The use of the software was explained. Subjects were also given notebooks and encouraged to write down any calculations that they wished to make, along with their decision for each round. These notebooks were used to inform the specification of the regression models tested below.

At the end of the session, the subject handed in their notebook and arranged an appointment for the next session. To refamiliarize subjects with the task, the instructor began each follow-up session with a summary of the objective of the task and the rules of the game.

RESULTS

Typical Behavior: Before turning to the statistical results, it is useful to consider the dynamics generated in typical trials. Figure 4 compares the final trials for two subjects, thus the results reflect subjects' decision processes after extensive experience. Overall performance for the two subjects is not statistically different from the grand mean for the full sample or from each other ($p>.7$). Both subjects faced the same random walk in exogenous sales, so their results are directly comparable. Subject 14 received the no-delay, no-gain condition, while subject 15 received a difficult condition with a long time delay and moderate positive gain ($\delta = 4, \gamma = +.3$). Subject 14, in the simple world with no feedback effects and no delays, is able to control the system quite well. Inventory remains within ±50 units of normal, and there is little tendency towards oscillation. Subject 15, in contrast, produces a large amplitude business cycle with a period of about 18 time units. Inventory deviates from normal by nearly ±500 units, ten times more than for subject 14. Over the course of the trial the cycle grows in amplitude – subject 15 seems to be losing control over time rather than learning to bring the system into equilibrium.

To see how feedback complexity leads subjects to create cycles, consider the beginning of the trial. By chance, exogenous sales fall, causing excess inventory to accumulate. Both subjects 14 and 15 respond with production cuts. For subject 14, the cut in production has an immediate effect on inventory and no induced effect on demand. The inventory imbalance is quickly eliminated, and subject 14 is then well positioned to deal with subsequent random changes in demand. In effect, subject 14 has 'reset' the system roughly to initial conditions (Edwards 1990); past decisions have little bearing on his current situation. Subject 15, however, finds first that his production cuts do not take effect immediately, causing still more excess inventory to accumulate, and leading him to still larger production cutbacks. Each unit cut from production, through the multiplier feedback, yields a .3 unit drop in endogenous demand. So even as he cuts production to eliminate excess inventory, demand falls, leading to additional inventory accumulation. By period 9 production has finally fallen below demand, and inventory begins to drop back towards the desired level. Failing to account for the time delay, subject 15 holds production at its depressed levels too long so that inventory falls below normal in period 12. As inventory plummets, the subject gradually boosts production, but inventory continues to fall as production increases lead to demand growth. Finally, by period 18, production exceeds demand and inventory begins to rise towards normal. Once again, however, the subject fails to account for the time delay, continuing to boost production until it exceeds demand by more than 140 units/period just when inventory reaches normal levels. Thus the cycle continues as excess inventory rapidly builds up, forcing another round of production cuts.

Note that the behavior of production and inventories is far from random. The peaks and troughs of production tend to lag the changes in demand, and that the amplitude of production is greater than the amplitude of demand. The amplification and phase lag are most pronounced for subject 15, but also are visible for subject 14. Amplification and phase lag are commonly observed in experimental tests of other stock management systems (Sterman 1989a, 1989b) and are a fundamental feature of business cycles in the actual economy (Mitchell 1927, Moore 1983).

Performance against benchmarks: To assess performance we calculate two benchmarks, the 'no-control' rule and the optimal rule. The no-control rule assumes subjects minimize cognitive effort and adjustment costs by making no changes in production; Δ_t is set identically to zero for all t. The optimal solution is somewhat more complex. The task system is linear, with quadratic costs. The optimal controller in such 'linear-quadratic' systems is a 'full state variable feedback' rule in which the optimal rule is a weighted sum of all state variables in the system (D'Azzo and Houpis 1981). The state variables are inventory, sales, current production, and the supply line of future production for $t = t+i$, $i = 1,...,\delta$. In equilibrium inventory is zero (the desired level) and production = sales. The rule strives to bring the system from its current state into equilibrium at the rate that optimally balances inventory costs and adjustment costs. It is important to note that the optimal rule does not utilize any information not presented to the subjects. The optimal weights are given by the solution to the Riccati equation (see e.g. D'Azzo and Houpis 1981).[2] Naturally, the optimal weights differ in the different delay and gain conditions, as do optimal costs.

Note that applying the optimal rule minimizes total *expected* costs. The optimal decision rule is not contingent on any particular realization of the random walk. The optimal rule correctly assumes exogenous sales follow a random walk so the best forecast of future sales is current sales. Subjects, however, may have guessed at the future values of exogenous sales. Being lucky and correctly anticipating the next values for exogenous sales can allow a subject to outperform the optimal rule occasionally (though not on average). Even the no-control rule, for instance, achieves lower costs than the optimal rule in five of the fifteen random walks in one experimental condition.

The benchmarks were calculated by running two sets of simulated subjects through the full design faced by the human subjects, yielding 225 'no-control' trials and 225 optimal trials. In figure 5 we plot subjects' average costs by treatment condition with the two benchmarks, ranking the experimental conditions from lowest to highest optimal costs. Because optimal costs in the easy and hard conditions differ by many orders of magnitude we plot the $\log_2$ of costs.

Comparison to 'No-control': Subjects' mean $\log_2$ cost was 18.28 ($318,000) and ranged from 12.46 ($5600) to 29.74 ($897 million). As expected, costs are highest in the most difficult conditions. Because raw performance confounds changes in decision behavior with changes in objective task difficulty, we turn now to comparison of subject costs to the benchmarks. Individual performance ranged from more than one thousand times better than no-control to more than 700 times worse. Overall, subjects achieve costs about 15 times lower than no-control.

Subject performance relative to the no-control strategy varies dramatically across treatment conditions (table 1; figure 6). The main effects of delay, gain and their interaction are significant at p = 000+, .000+, and .026, respectively. In the three most favorable conditions ($\delta = 0$; $\gamma = -0.6$, -0.3, and 0) the majority of subjects outperform the no-control rule on average by more than a factor of 100. However, the no-control rule outperforms the subjects on average in the two least favorable conditions ($\delta = 2$ and 4; $\gamma = +0.6$). In these conditions subjects would have been better off not making any changes in production at all – their efforts to control the system were counterproductive. While the best subject outperforms the no-control rule in the most favorable condition ($\delta = 0$; $\gamma = -0.6$) by more than 1000 times, the best subject in the least favorable condition ($\delta = 4$; $\gamma = +0.6$) outperforms no-control barely more than four times. Relative to the no-control rule, performance deteriorates with increasing delay and performance decreases with increasing gain. In addition, delay and gain interact, with long delay times and high positive gains causing a particularly steep decrease in performance relative to the no-control benchmark.

The practice/sequence effect is highly significant as well. Recall that the practice/sequence effect is attributable to two parts: learning and differences among the random walks. However, it appears that performance in the first two trials contains greater variance and more poor performers. The no-control rule outperforms many subjects in their first trial. Performance in subsequent trials is clearly improved, though after about three trials there is only slight further improvement. Performance relative to the no control rule varies considerably between subjects, as expected. Subjects who do poorly in the easier trials do especially badly in the most difficult cases.

Comparison to Optimal: While subjects generally outperform the no-control rule, they perform quite poorly relative to optimal. Subject costs are 4.35 times greater than the corresponding optimal performance across all conditions. Most interesting, however, is the relationship between subject costs and optimal costs. Subject performance remains roughly parallel to but consistently higher than optimal as difficulty increases (figure 5). Indeed, the ANOVA (table 2) shows no significant effect of either treatment or their interaction, suggesting subjects did adjust their decision making heuristic as the delay and gain conditions changed.

As expected, subjects' performances relative to optimal differ, and there is a significant practice/sequence effect. Scores are highest in the first two trials and continue to improve slightly thereafter, consistent with the raw performance results.

Decision Effort: We analyzed the time spent per decision for both the initial decision in each new condition and later decisions. In all cases a significant practice effect is expected, as are intersubjective differences. Predictions about the dependence of time spent on delay and gain differ according to the different hypotheses about subject decision behavior. The extreme MOF hypothesis suggests people are insensitive to delays and gain, focusing only on the current value of the stock. Thus time spent, either in the orientation phase before a new trial or during a trial, should not depend on either treatment condition. Both a more moderate MOF hypothesis and the rational model suggest people try to account for delays and side effects. The time spent in orientation before the first decision should thus increase with increasing delay and gain compared to the no-delay, zero-gain condition. The complexity of the solution to the optimal rule rises roughly with the square of the number of state variables but does not depend on the gain of the side-effect loop (except when γ=0). If subjects approximated the optimal strategy their initial decision times should show similar differences. Likewise, the optimal rule is a weighted average of all state variables in the system. The number of states is independent of the gain, and depends linearly on the delay condition. Thus subjects using a full state variable decision rule should exhibit longer subsequent decision times in the delay conditions, but decision time should not vary with the gain. Under the misperception of feedback hypothesis, subjects attempt to account for the delays and side effects, but do so imperfectly, for example by accounting only for the most recent production decision rather than all stages of the supply pipeline. Here initial decision times should depend weakly on the delay and gain conditions. Later decision times might also depend on the treatments as the larger gains might induce greater attention to the side-effect loop, and longer delays might induce greater effort to account for the supply line.

Table 3 shows the ANOVA results for initial and subsequent decision times. As expected, there are significant differences between subjects in the time taken, and there is a substantial practice effect, with decision times falling rapidly in the early trials.

There is no apparent relationship between the treatment conditions and the time spent on the initial decision in each condition. Likewise, the main effects of delay and gain on subsequent decision times are insignificant. Though the interaction of gain and delay is significant at about the .01 level, there is no clear pattern to the differences among conditions. Overall, each subject spent the same amount of time making decisions, both at the start of each new trial and subsequently, in all treatment conditions, despite large variations in task difficulty and number of available cues.

The decision timing data support the MOF hypothesis that subjects are insensitive to and/or unable to account for the feedback structure of the task, and are not consistent with the rational model, which predicts longer decision times in the more complex conditions. The rational model also predicts that the time taken at the beginning of a new trial should be longer in more complex conditions, yet there is no evidence this is so. Dynamic complexity does not lead to increased effort. Further, increased effort does not lead to greater performance: though subjects differ greatly in the time they spend in decision making, regressions show no significant relationship between performance relative to the benchmarks and either initial or subsequent decision times.

The fact that effort does not improve performance may be caused by several factors. Mental processing speed may vary among subjects independent of competence in the task. Extra time may be spent in irrelevant thought, rest, or distraction. Those who understand the task poorly may compensate by spending additional time. Finally, people's understanding of dynamic complexity and its implications may be so poor that additional effort is not helpful. To explore these issues we turn now to models of the subjects' decision rules.

MODELING SUBJECTS' DECISION RULES

To gain deeper insight into the nature of the decision process we fit various linear models to the subjects' decisions. The specification of these rules was based on prior experimental work, the feedback structure of the task, and a detailed analysis of the subjects' notebooks. Subjects received a notebook in which they wrote their decisions in the upper right hand corner, while other calculations could be made on the rest of the page. Analysis of the notebooks consisted of reviewing subjects' round by round decisions, while relating the written calculations and comments to each decision. The analysis was done for half of the subjects, selected randomly.

The first three games were examined to see how subjects initially approached the task. These initial games were compared with selected later games with particular emphasis on long delay and positive gain, in order to determine how strategies formulated early on held up over time and under the difficult conditions. All games with long delays and high gains were investigated.
The analysis of each game considered a variety of factors including average time overall and per game, including timing differences within a game, the number of calculations recorded, attention to/proper calculation of the effects of feedback gain, attention to/proper calculation of production delay, over/under emphasis of inventory, and over/under emphasis of the change in inventory.

While such notebook analysis is a valuable addition to the researcher's tool kit of process methods (Carroll and Johnson 1990), it has limitations. In particular, writing down calculations and performing them are two different processes. Just because nothing is in the notebook does not necessarily mean that subjects did not perform calculations in their heads.

The analysis revealed three major levels of sophistication:

1) attention to inventory only: subjects respond only to the inventory discrepancy;

2) attention to inventory and change in inventory: subjects understand the relationships between production and sales and respond to the rate of change in inventory as well as the current stock;

3) attention given to inventory and *expected* change in inventory: subjects attempt to account for the supply-line.

While most subjects exhibited at least some awareness of the importance of the change in production and the supply line, the notebooks revealed that subjects do not give these factors sufficient consideration in high delay and high gain conditions. Some of the subjects' calculations showed they estimated future change in inventory by ignoring the side effect loop entirely, or miscalculating its impact. While not totally ignoring the supply line, it appears that subjects paid less than full attention to future changes in inventory. Many exhibited a marked tendency towards undercontrol, reasoning, as subject 12 wrote, "Gradual changes seem the best since independent is so random that big changes cost a lot and have a high probability of over- or under-shoot." Subject 13 likewise wrote "sales independent unpredictable, cut cost by keeping production constant."

In addition, the notebook analysis revealed a number of interesting strategies subjects used to explore the system. Many subjects adopted an 'act and wait' strategy. After changing production these subjects would often enter 0 for the change in production in the next few rounds, apparently to wait and see what the effects of the prior change would be before exerting additional control. Some subjects made extensive calculations initially, examining the effects of the delay and the side-effect loop, but abandoned such sophisticated decision making when their decisions created instability, reverting to an inventory-only heuristic. Others, faced with oscillations created by their failure to account properly for the delay and side effect, reverted to the no-control strategy, apparently hoping the system would settle down so they could resume control. Subject 4, for example, made extensive calculations of inventory changes in the early games, though his calculations ignored the side effect loop. In the difficult condition $\delta = 4$, $\gamma = +.6$, such behavior yields costly oscillations, prompting the subject to write "I can't believe this is happening!" The subject then ceased to write any calculations and seemed to revert to a strategy focused only on the current inventory discrepancy.

One subject cleverly explored the dynamics of the system by implementing production changes of .01 units. These small changes allowed the subject to explore the system's response at low cost, and are reminiscent of the engineering technique of 'small signal analysis.' However the subject soon abandoned the technique in favor of an inventory-only strategy. Some exhibited the illusion of control by hoping the exogenous component of sales would change so as to eliminate their inventory discrepancy without the need for them to make costly changes in production. Subject 2, in his first trial, facing an inventory surplus, wrote "Hope that independent sales will go up." His production cutback is much less than the amount needed to eliminate the excess inventory. The environment, however, disappoints him; independent sales drop from 927 to 913, and the subject writes "uh oh! Sales went down."

Drawing on the notebook analysis and hypotheses discussed above, four models of increasing sophistication were constructed. The rules consist of both the specification of the cue-combination policy and restrictions on the signs of the cue weights. The models range from the simplest single-cue rule to the optimal rule.

Model 1: $\Delta_t = \alpha_1 I_t; \alpha_1 < 0.$ (9)

Model 1 represents the extreme misperception of feedback hypothesis in which subjects respond only to inventory.[3] The coefficient α_1 must be negative so that positive inventories lead to production cutbacks and vice versa, forming a negative feedback loop. The notebook analysis showed all subjects used rules more sophisticated than Model 1, but we include it here because it is the control rule demanding the least cognitive effort and serves as a benchmark for comparison to other rules.

Model 2: $\Delta_t = \alpha_1 I_t + \alpha_2 (P_t - S_t); \alpha_1 < 0, \alpha_2 < 0.$ (10)

Model 2 assumes that subjects consider both inventory and the rate of change in inventory, given by Production - Sales. The coefficient α_2 must be negative. For example, when production exceeds sales inventory will increase, leading to a cutback in production. From a control-theoretic viewpoint, attending to the rate of change in inventory adds 'derivative control', a control strategy that opposes trends in the stock and is thus stabilizing (Ogata 1970). The notebook analysis revealed that some, although not the majority, of the subjects might have followed this rule. Note that although the rule is sensitive to the flows that alter the stock it does not account at all for the supply line of future production (if any).

Model 3: $\Delta_t = \alpha_1 I_t + \alpha_2 (P_{t+\delta} - S_t); \alpha_1 < 0, \alpha_2 < 0.$ (11)

Model 3 is similar to Model 2 except that subjects are assumed to account for the supply line of pending production by comparing future production to current sales in the derivative control term. Subjects know the value of production at time $t+\delta$ because $P_{t+\delta} = P^*_t$, and these values are displayed on the screen. However, Model 3 does not account for all stages of production in the supply line (P_{t+i}, $i = 1,...,\delta$) even though these values are helpful and are also displayed. Instead of comparing only the most recent production start decision to sales, subjects should consider the quantity in each stage of the supply line and compare these to expected sales in each period. Model

3 is thus not rational from a control theoretic viewpoint. However, model 3 does economize on the number of cues considered and the amount of cognitive effort required to process them (Kleinmuntz 1993). The model requires subjects to consider only one stage of the production delay instead of the three or five (δ+1) stages in the delay conditions.

$$\text{Model 4: } \Delta_t = \alpha_1 I_t + \alpha_2 S_t + \sum_{i=0}^{\delta} \beta_i P_{t+i}; \;\; \alpha_1 < 0, \alpha_2 > 0, \Sigma\beta_i = -\alpha_2 \tag{12}$$

Model 4 is the optimal decision rule, the full state variable feedback controller. Here inventory control is supplemented by comparison of current production and all stages of the supply line to expected sales from the current value through δ periods ahead. Sales increases deplete inventory, necessitating higher production, thus $\alpha_2 > 0$. The larger the current rate of production or the pending production in the supply line, the greater inventory will be, thus $\Sigma\beta_i < 0$. Furthermore, the production weights must sum to the (negative of) the weight on current sales to ensure that in equilibrium the change in production $\Delta = 0$ when $I = 0$ and sales = production. Although the notebook analysis did not support the full information model, it is important as a test of the sophistication of the subjects' decision making, and as a measure of the extent to which subjects attended to the supply line.

All four models were estimated by OLS for all 225 conditions. Adjusted R^2, the significance of the estimated coefficients, and the number of coefficients with incorrect signs were used as indicators of the adequacy of each in capturing the subjects' decisions. For instance, it makes no sense at all to increase production in response to a positive inventory. Such a policy would create a reinforcing feedback whereby any initial inventory discrepancy is amplified, leading to unbounded costs. Thus a positive weight on inventory would be judged an artifactual sign reversal and evidence the model was misspecified. Any model with one or more incorrect signs for the coefficients was classified as a case of sign reversal. Table 4 shows the average adjusted R^2 and number of sign reversals for each model.

Model 1 is clearly dominated by the others, with a mean $\overline{R^2}$ of just 26% and 24 sign reversals, despite having only one parameter. Adding derivative control as in Model 2 raises $\overline{R}^2$ to 51% but also reveals 59 cases with at least one sign reversal, more than one quarter of the conditions. Model 3 dominates Model 2, with a higher $\overline{R}^2$ and less than half the incidence of incorrect signs for the coefficients. Fully 14 of the 24 cases with incorrect signs occur in the conditions with the longest delay and the highest positive gains, suggesting subjects abandoned model 3 in these difficult conditions.

The full information Model 4 increases $\overline{R}^2$ by .04 compared to model 3. However, 136 of the regressions (60%) have at least one coefficient with the wrong sign. Nearly all the sign reversals occur in the delay conditions (model 4 reduces to model 2 in the no-delay conditions). Only 2 of 75 cases had the correct signs in the conditions where $\delta=4$. While it is true that model 4 includes more coefficients than model 3, so the chance of an incorrect sign is greater, the high incidence of incorrect cue weights strongly suggests the subjects were not using the full information model. Furthermore, only 11% of the estimated weights for intermediate stages of the supply line were significantly different from zero (at the 5% level), indicating subjects did not attend to all stages of the supply line, as is optimal, instead ignoring production they had initiated but not yet received, as predicted by the MOF hypothesis.

Besides the four models discussed in this section, seven models ranging in sophistication between models 3 and 4 were also tested. These models consisted of adding various states intermediate between model 3 and the complete state model in order to test whether subjects' behavior might fall in sophistication somewhere between model 3 and the full-state model. All of the alternative models fell in the range $.58 \le \overline{R}^2 \le .62$; also, all of these models produced sign reversals substantially higher than model 3.

Thus while model 4 and related models in which intermediate supply line states are considered explain slightly more of the total variance in subject's decisions, we judge them to be inferior to model 3 as representations of cue utilization and weighting. Analysis of subject's notebooks also showed that the majority learned to attend to future changes in inventory but did not learn to attend to all stages of the supply line, thus supporting the statistical results favoring model 3.

We also analyzed subject and treatment differences in the quality of model fit to examine the consistency of use of Model 3 across subjects, trials and treatments. Indeed, the goodness of fit of Model 3 is significantly different across experimental conditions (p = .000+). Fit is best in the no-delay, negative gain conditions, and deteriorates as delay and gain increase. In the most difficult conditions ($\delta = 4, \gamma = +.6$) the average $\overline{R}^2$ is just 27%, compared to 83% in the easiest condition ($\delta = 0, \gamma = -.6$). In the difficult conditions the system is intrinsically less stable and costs are substantially higher. The higher variances of the cues might cause greater inconsistency in the sense of inadvertent error caused by rounding, carelessness, or decreased effort caused by the perception of poor performance. Alternatively, reduced consistency may reflect deliberate experimentation and search for better policies. The notebooks suggest both occur, though the incidence of explicit experimentation is low. The quality of fit also differs across subjects. Perhaps surprisingly, the practice/sequence effect is not significant. Subjects apparently did not become more consistent with experience.

Sources of underperformance

Underperformance can arise in two principal ways: (1) Subjects employ the right model but apply it inconsistently; (2) Subjects employ a poor model. To partition the total shortfall of performance relative to optimal into these two components we perform a bootstrap analysis 'running' the experiment again using the estimated Model 3 decision rules for each of the 225 cells of the design. The performance of the estimated rules compared to that of the subjects measures the effect of subject inconsistency. The performance of the estimated rules relative to the optimal rule measures the effect of using the wrong model. Figure 7 shows that the bootstrap simulations yield a substantial improvement, with simulated costs about 3 times greater than optimal, compared to 4.35 times greater for the subjects. The improvement of about 30% is large relative to the typical bootstrap effect in static tasks (Camerer 1981, Dawes and Corrigan 1974). More detailed examination shows the bootstrap effect varies with the experimental conditions. The bootstrap effect is most pronounced in the no-delay condition. With long delays and high positive gain many subjects outperform the bootstrap model.

Despite the substantial bootstrap effect, three-quarters of the underperformance relative to optimal is due to subject's use of the wrong model, particularly underweighting of inventory and future production and failure to attend to all stages of the supply line. The misperceptions of feedback dominate the effects of inconsistency. Subjects are not able to account adequately for the time delay or side-effect feedback. Because the existence and strength of these structural features are fully revealed, and full information about these effects is available on the information display, the subjects' failure to respond appropriately must reflect more fundamental limitations on the processing of information about the elements of feedback complexity.

Dependence of Cue Weights on Experimental Conditions

We now turn to dependence of cues weights on the experimental conditions. How far from optimal are the cue weights used by the subjects? How do the cue weights change as delay and gain change? Appropriate changes in cue weights across experimental conditions would provide evidence that subjects adjusted their decision strategies to account for the feedback structure of the system. Constant cue weights across experimental conditions would support the extreme misperceptions of feedback hypothesis: subjects would be completely insensitive to the feedback structure of the task. Partial adjustment of cue weights would support a more moderate MOF hypothesis. Figure 8 shows 15 graphs showing estimated cue weights for inventory (vertical axis) and future change in inventory (horizontal axis). The rows show the 3 delay conditions and the columns show the 5 gain conditions. Also shown are the optimal weights for model 3.[4]

Inspection of figure 8 reveals that nearly all of the subjects' weights lie to the lower left of optimum, corroborating the earlier result of significant undercontrol, or insufficient responsiveness to inventory and flow imbalances. The degree of undercontrol is not constant across treatment conditions:

1. Under increasing complexity, the optimal weight for inventory stays roughly the same, but subjects' weights for inventory decrease dramatically.

2. Under increasing complexity, the optimal rule places greater and greater weight on future change in inventory, but subjects' weights remain about same.

To verify these findings, we calculate for each experimental condition the Euclidean distance, D, between the optimal weights and each subject's weights for inventory and future change in inventory, given by

$$D_{\delta,\gamma,i} = [(\alpha_{1,\delta,\gamma,i} - \alpha^*_{1,\delta,\gamma})^2 + (\alpha_{2,\delta,\gamma,i} - \alpha^*_{2,\delta,\gamma})^2]^{1/2}; \quad i = 1,...,15. \tag{13}$$

The ANOVA analysis of D (table 6) shows that the distance between subjects' weights and optimal weights varies significantly with delay, gain, and their interaction, all at the p=.000+ level. The deviation from the optimal weights increases as feedback complexity increases.[5] In addition, subject differences are significant, but the practice/sequence effect is not. As the feedback complexity of the task increases, subjects increasingly undercontrol the system. In particular, the optimal weight on inventory remains approximately constant across experimental conditions, but subjects responsiveness to inventory falls as delay lengthens and gain increases. On average, subjects assign a weight of -0.15 to inventory. The optimal rule assigns a weight to inventory of about -.39. Under zero delay, the average weight is approximately -.25; when δ=2, the average weight drops to approximately -.14 and shows greater variability across gain conditions. When δ=4, the average weight drops to approximately -.05 and decreases as gain increases – subjects move towards the 'no-control' rule as the task becomes more difficult.

The optimal weight on the difference between future production and sales increases substantially as the gain increases, in order to compensate for the positive loop created by the side effect of production on sales. Yet the subjects are not more responsive to the anticipated change in inventory as gain increases. The weights associated with the future change in inventory stay roughly the same across conditions except for the $\delta = 4$ condition. Subjects do not adjust their cue weights appropriately as the dynamic complexity of the system changes. The insignificant practice/sequence effect suggests as well that they fail to learn to adjust their decision weights over time as complexity changes, despite fifteen trials each of thirty two decision rounds.

Given the large and significant discrepancies between optimal and subject weights across treatments, it is surprising that there are not greater increases in the ratio of subjects' costs to optimum as feedback complexity grows. To explore the insensitivity of costs to cue weights, we mapped the cost surface of the task in each of the fifteen conditions as functions of the cue weights for model 3 (figure 9). The black dot in the center of each graph represents minimum costs. Costs are normalized so that each contour line marks a 50% increase in costs from optimal. The cost surfaces form valleys with steep sides and broad bottoms. There is a relatively large region of low costs around the optimal cue weights, indicating the task has a flat optimum, as observed in a number of prior dynamic decision making studies (e.g. Rapoport 1975). Furthermore, as delay and gain increase, the floor of the valley narrows and stretches out toward the origin: as feedback complexity increases, the optimal strategy converges to the 'no-control' rule. The flat cost surface of the task and shift towards no-control of the optimum cue weights as complexity grows thus explains the weak dependence of subject costs relative to optimal as difficulty increases even though subjects do not adjust their decision rules appropriately as feedback complexity grows.

DISCUSSION

There are two competing explanations for the significant differences in subject weights relative to optimal across delay and gain conditions. One could argue that subjects understood the structure of the task well enough to conclude that the task has a flat optimum, and that they understood that the valley of low costs around that optimum point moves towards the 'no-control' policy as feedback complexity increases. Thus the insufficient attention paid to inventory and supply line cues represents a rational, or at least reasonable, tradeoff between performance and effort, so that performance relative to optimal remains relatively constant across treatments. However, reducing the weighting of inventory and the supply line does not reduce the need to monitor those cues, nor does it reduce the need to combine the cues to make a decision. Thus the cognitive effort saved by underweighting seems modest at best, while the average shortfall of costs from optimal was large (more than a factor of four). Further, deciding what cue weights to use as a deliberate strategy based on understanding of task structure in each gain and delay condition would require subjects to spend time at the start of each trial determining how the gain and delay conditions of that trial affected the cost surface. The time required to calculate this surface should be longer in conditions with long delays, since the system involves more state variables (the complexity of finding the cost surface rises roughly with the square of the number of state variables). Yet the deliberation time data showed no significant dependence on delay or gain, for either the initial or subsequent decisions in each conditions. Furthermore, effort (as measured by mean decision time) was uncorrelated with performance. Finally, the notebook analysis provides virtually no support for the hypothesis that subjects chose their decision rules by explicit consideration of the delay and gain conditions.

It is more likely that subjects do not understand the response surface of the task or the dependence of optimal cue weights on delay and gain conditions. By this interpretation under-weighting increases with complexity because subjects suffer from two types of misperceptions of feedback. They are insufficiently aware of the feedback structure of the system, and they are not able to infer the consequences of the feedbacks and time delays. As delays lengthen and gain increases the system becomes less stable and the optimal policy requires more aggressive response to the supply line and the anticipated change in production. Yet subjects, unable to infer properly how the gain and delay conditions affect system behavior, fail to increase their responsiveness to the anticipated change in production. Subjects reduce their responsiveness to inventory as complexity increases because they find the system to be less stable – failing to understand how their control actions affect the system, they reduce the control they exert, moving to a hands-off strategy. Much of the data are supportive of the MOF hypothesis. The subjects tend to produce costly, persistent oscillations with the characteristic amplitude and phase relations observed in prior work and predicted from simple models in which the feedback structure and time delays are ignored. Regression analysis of subject decisions showed most were unresponsive to the time delays, ignoring the supply line of pending production. The deviation of subject cue weights from optimal grows larger as the dynamic complexity of the system grows. Finally, the notebook analysis suggests most subjects did not attempt, at least for very long, to account for the delay and feedback loop structure of the task, tending instead to revert to a simple model focusing on inventory and perhaps the expected change in inventory. In a sense, subjects were lucky: though performance was not very good overall, their misperceptions of the feedback environment and its consequences were not punished more severely in the difficult conditions because the task has a relatively flat cost surface.

Conditions in the experiment favor high performance compared to many real-life dynamic decision tasks. The subjects, though students, rank much higher than the average person in relevant knowledge of mathematics, computers, and feedback control theory. There were monetary incentives for performing well. Explanation of and training in the task were provided. Subjects could take all the time they wished to make their decisions. Full information was presented to the subjects; there were no hidden states. The number of cues to attend to was small. The feedback structure of the task, including full disclosure of the treatment conditions, was displayed at all times. The task, except for the exogenous component of sales, was completely deterministic – delay length, feedback strength, outcome feedback, and subject decisions were not subject to any random variation or measurement error. Outcome feedback was immediate, perfect, and complete.

Yet overall performance is quite poor compared to optimal. In the most difficult conditions, subjects are outperformed on average by the 'no-control' rule. Analysis of subjects' decisions revealed that they tended to focus on inventory (the stock to be managed) and the difference between expected production and sales, rather than considering full information. As found in prior experiments, subjects did not attend to the supply line, even though information about its contents was just as prominent in the information display as information about inventory or sales. Side effect feedbacks also proved difficult to manage. In the presence of a positive side-effect loop, where the side effect undercuts the impact of the subjects' control efforts, subjects undercontrolled the system dramatically. Most generate costly oscillations as they failed to account properly for the impact of the side effect feedback.

The results strongly support the misperceptions of feedback hypothesis. Because the subjects had full information, training, incentives, and extensive opportunities for experience it is not plausible to attribute their poor performance in the complex conditions to inadequate information, ambiguity of outcome feedback, inexperience, or lack of effort. Though these issues are indeed grave impediments to effective performance in many real life settings, these results suggest the source of the problem is more fundamental.

We suggest the mental models subjects bring to bear in complex tasks are dynamically deficient. Subjects were unable to account well for delays and feedback effects because (1) people's mental models of control tasks are highly simplified, tending to exclude side effects, feedback processes, delays, and other elements of dynamic complexity; and (2) even when these elements are known, people's ability to infer correctly the behavior of such complex feedback systems is poor. The former misperception can be overcome with training in the principles of feedback systems and dynamics. The latter is a fundamental bound on human rationality – our cognitive capabilities do not include the ability to solve systems of high-order nonlinear differential equations intuitively. Fortunately, computer simulation makes the task of solving these systems trivial. Currently available software allows people from grade school to chief executive to build and simulate dynamic systems of arbitrary complexity on their personal computers (Richmond 1993, Diehl 1992, Eberlein and Peterson 1992, Morecroft and Sterman 1992). Nevertheless, the art of model building remains difficult. Exploring the nature of training and decision aids to overcome the misperceptions of feedback and testing their effectiveness in the field – in firms, markets, and other real dynamic decision making contexts – is the next frontier for research in the psychology of dynamic decision making.

NOTES

1 Technically, the delay introduces phase lag that can cause the eigenvalues of the linearized system to become complex conjugates. Such systems are oscillatory. In the general case, the system may be damped, stable but underdamped, or unstable, depending on the gain of the control loop (the response to discrepancies between the stock and its goal) and the length of the time delay (Forrester 1961, Ogata 1970).

2 The weights were calculated numerically using Matrix X as implemented on MIT's Athena system.

3 All four models were tested both with and without a constant term. The constant is predicted to be zero; nonzero values would yield nonzero inventory costs in equilibrium. In all cases the estimated constants were not significanlty different from zero, as expected, so the regressions reported here were run with the constant suppressed.

4 The optimal weights for model 3 were derived numerically (Diehl 1992).

5 ANOVA analysis of each weight separately confirms these results (Diehl 1992).

REFERENCES

Bakken, B. (1993) Learning and Transfer of Understanding in Dynamic Decision Environments. Ph.D. Dissertation, MIT Sloan School of Management, Cambridge, MA 02142.

Brehmer, B. (1990) Strategies in Real Time, Dynamic Decision Making, in Hogarth, R. (ed.) Insights in Decision Making. Chicago: University of Chicago Press, 262-279.

Brehmer, B. (1992) Dynamic Decision Making: Human Control of Complex Systems. *Acta Psychologica* 81, 211-241.

Camerer, C., "General Conditions for the Success of Bootstrapping Models," *Organizational Behavior and Human Performance*, 27 (1981), 411-422.

Carroll, J. S., and Johnson, E. J. (1990). Decision Research. Newbury Park: Sage Publications.

Dawes, R. and Corrigan, B. (1974) Linear models in decision making. *Psychological Bulletin. 81, 95-106.*

D'Azzo, J. J., and Houpis, C. H. (1981). Linear Control System Analysis and Design: Conventional and Modern . New York: McGraw-Hill Book Company.

Diehl, E. (1992) Effects of Feedback Structure on Dynamic Decision Making. Ph.D.. Dissertation, MIT Sloan School of Management.

Diehl, E. W. (1992). Participatory simulations for managers: The design philosophy behind MicroWorld Creator. European Journal of Operations Research, 59(1), 210-215..

Diehl, E. W. (1989). A study on human control in stock-adjustment tasks. In P. M. Milling, & E. O. K. Zahn (Ed.), Computer-based management of complex systems (pp. 205-212). Berlin: Springer.

Eberlein, R. and Peterson, D (1992) Understanding Models with VENSIM. European Journal of Operations Research, 59(1), 216-219.

Edwards, W. (1990) Unfinished Tasks: A Research Agenda for Behavioral Decision Theory," in Hogarth, R. (ed.) Insights in Decision Making: A Tribute to Hillel J. Einhorn. Chicago: University of Chicago, 44-65.

Forrester, J. W., 1961. *Industrial Dynamics*. Cambridge: MIT Press.

Funke, J. (1991) Solving Complex Problems: Exploration and Control of Complex Systems, in R. Sternberg and P. Frensch (eds.), *Complex Problem Solving: Principles and Mechanisms*. Hillsdale, NJ: Lawrence Erlbaum Associates.

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

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

Kampmann, C. and J. Sterman (1992) Do Markets Mitigate Misperceptions of Feedback in Dynamic Tasks? Proceedings of the 1992 International System Dynamics Conference, Utrecht, The Netherlands.

Kleinmuntz, D. (1993) Information processing and the misperception of feedback in dynamic decision making. *System Dynamics Review*. forthcoming.

MacKinnon, A. J., and Wearing, A. J. (1985). Systems analysis and dynamic decision making. Acta Psychologica, 58, 159-172.

Mitchell, W.C. 1927. *Business Cycles: The Problem and Its Setting*, New York: National Bureau of Economic Research.

Moore, G. (1983) *Business Cycles, Inflation, and Forecasting*. Cambridge, MA: Ballinger/National Bureau of Economic Research.

Morecroft, J. and J. Sterman (eds.) (1992) Modeling for Learning. Special issue of *European Journal of Operational Research*. 59, 1 (May).

Ogata, K. (1970) *Modern Control Engineering*. Englewood Cliffs, NJ: Prentice Hall.

Paich, M., & Sterman, J. D. (1993). Boom, Bust, and Failures to Learn in Experimental Markets. *Management Science*., forthcoming.

Rapoport, A. (1975). Research paradigms for studying dynamic decision behavior. In D. Wendt, & C. A. J. Vlek (Ed.), *Utility, Probability, and Human Decision Making* . Dordrecht, Holland: Reidel, 349-375.

Richardson, G. (1991) *Feedback Thought in Social Science and Systems Theory*. Philadelphia: Univ. of Penn. Press.

Richmond, B. (1993) Systems thinking: critical thinking skills for the '90s and beyond. *System Dynamics Review*, 9(2), 113-134.

Smith, V., Suchanek, G., and A. Williams (1988) Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets, *Econometrica*, 56(5), 1119-1152.

Sterman, J. D. (1989a) Misperceptions of Feedback in Dynamic Decision Making. *Organizational Behavior and Human Decision Processes*, 43(3), 301-335.

Sterman, J. D. (1989b) Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science*, 35(3), 321-339.

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. (1989). Deterministic Chaos in an Experimental Economic System. *Journal of Economic Behavior and Organization*, 12, 1-28.

Sterman, J. D. (1987). Testing Behavioral Simulation Models by Direct Experiment. *Management Science*, 33(12), 1572-1592.

Figure 1. The generic stock management structure, as implemented in the experiment. Top: the stock and flow structure of the task, showing the two treatment variables, production delay, δ, and side-effect feedback gain, γ. Bottom: Subjects choose the change in production, Δ, by considering their inventory position relative to the target, and possibly accounting for the supply line of future production and the expected future value of sales.

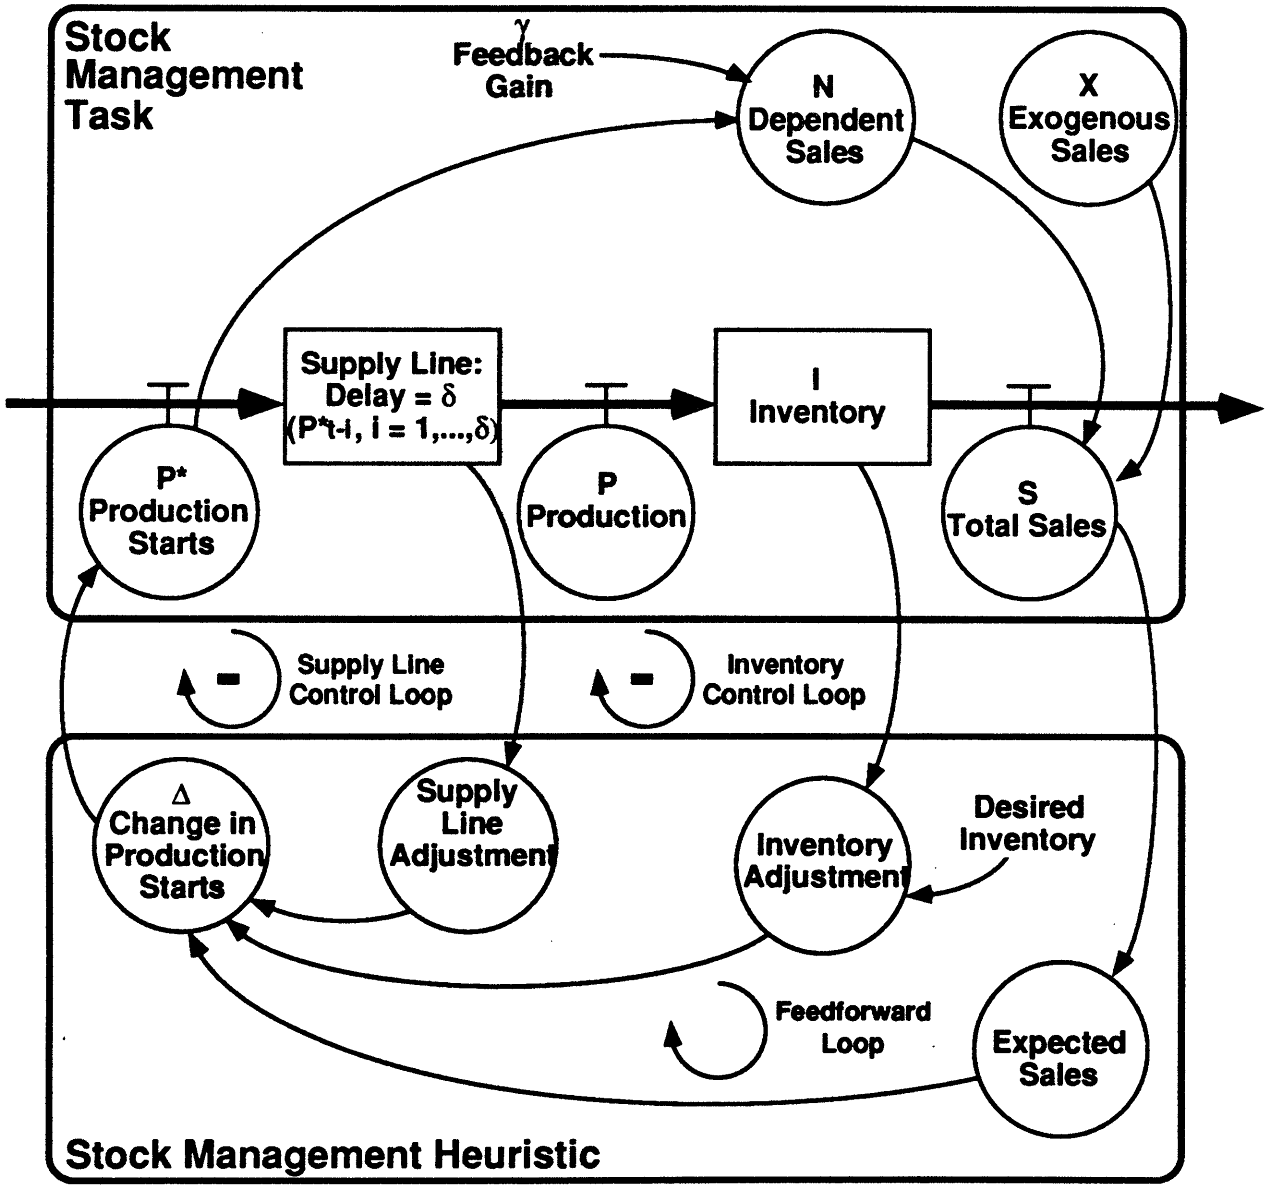

Figure 2. Subject/sequence table and Latin-square with 3 levels of production delay and 5 levels of feedback gain.

Subject number \ Trial number	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	I	B	N	G	A	F	C	M	H	J	D	L	O	K	E
2	K	D	G	F	C	E	O	J	A	M	B	H	I	N	L
3	L	A	M	O	D	J	G	B	E	N	F	K	C	I	H
4	B	F	J	H	N	D	A	K	I	E	L	M	G	C	O
5	N	O	A	I	M	H	J	D	C	L	G	B	F	E	K
6	C	H	E	M	B	N	K	A	F	O	J	I	L	G	D
7	O	M	I	C	J	L	H	G	D	F	K	N	E	B	A
8	J	E	C	D	L	O	I	H	M	K	N	G	A	F	B
9	F	L	K	J	E	B	N	C	O	H	I	A	M	D	G
10	M	I	H	N	G	K	L	O	B	C	E	J	D	A	F
11	G	C	B	E	O	M	D	L	N	A	H	F	K	J	I
12	D	G	L	K	F	C	E	I	J	B	A	O	N	H	M
13	A	N	D	L	H	I	M	F	K	G	C	E	B	O	J
14	E	J	F	A	K	G	B	N	L	I	O	D	H	M	C
15	H	K	O	B	I	A	F	E	G	D	M	C	J	L	N

Trial number

Subject per sequence table

Delay \ Gain	-0.6	-0.3	0	+0.3	+0.6
0	A	B	C	D	E
2	F	G	H	I	J
4	K	L	M	N	O

Latin Square design matrix

Figure 3. Sample computer screen display.

Week	59	60	61	62	63
Change in Production	0	0		Enter Decision	
Production	600	600	600	600	
Sales (dependent)	180	180			
Sales (independent)	420	423			
Sales (total)	0	603			
Change in Inventory	0	-3			
Inventory	0	-3			
Cost (Prod. Change)	0	0	= 2 * (0*0)		
Cost (Inventory)	0	9	= 1 * (-3*-3)		
Cost (Total)	0	9			
Accumulated Cost	0	9			

Conditions for current game:
- 10 production units cause 3 sales units.
- Production is delayed by 2 weeks.

Figure 4. Typical results.

Left: Subject 14, Final trial, Delay = 0, Gain = 0;

Right: Subject 15, Final Trial, Delay = 4, Gain = +.3.

Both subjects faced the same random walk in exogenous sales, so their results are directly comparable.

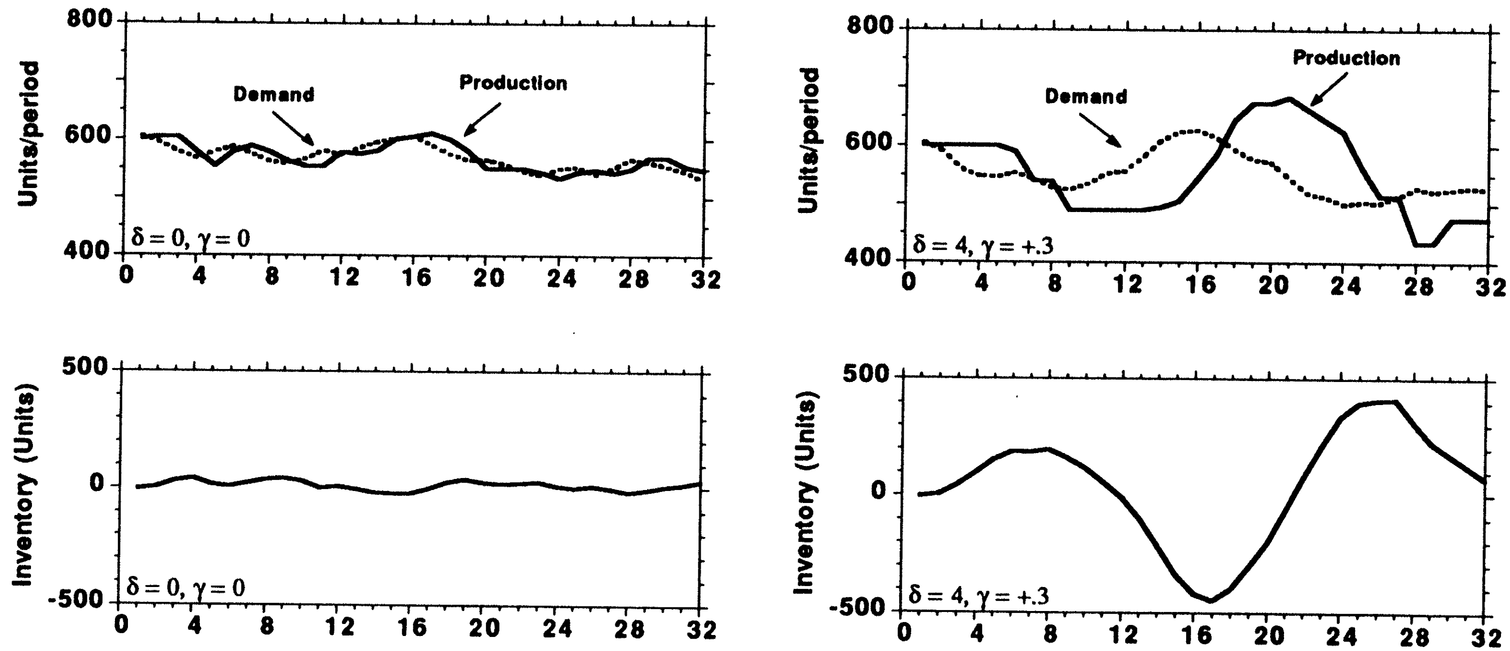

Figure 5 Average subject performance, log (base 2), compared to average of optimal rule and no-control rule. Experimental conditions are ranked from lowest optimal costs to highest.

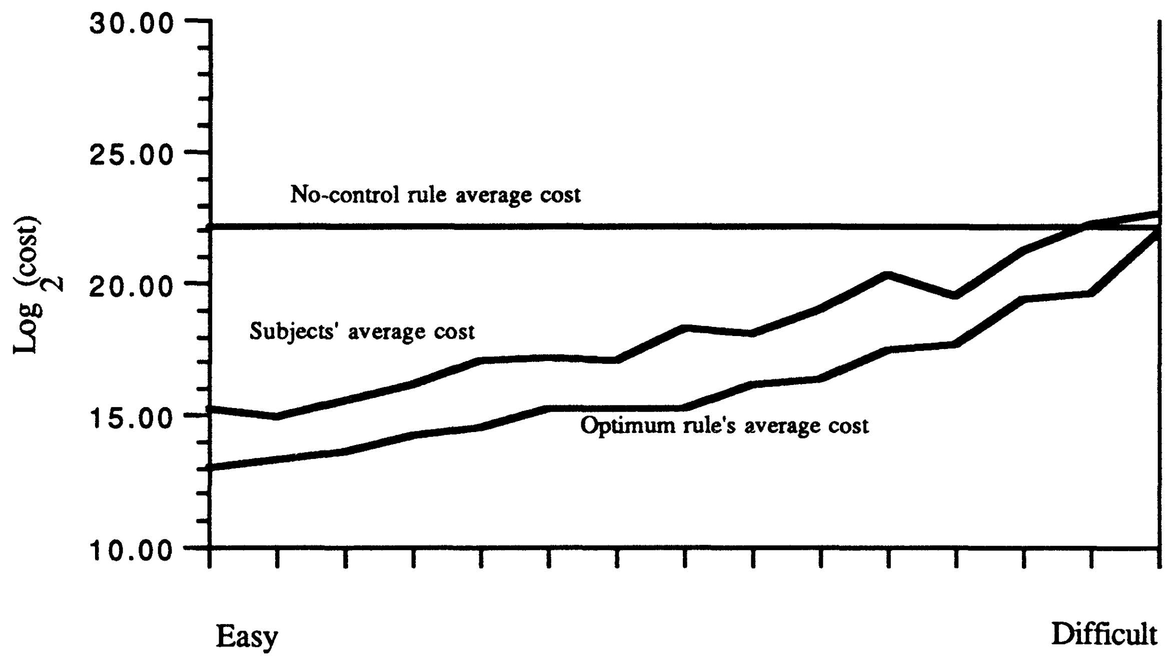

Figure 6 $\log_2$(Subject costs/no-control costs) by treatment condition. Negative (positive) numbers indicate the subjects outperformed (were outperformed by) the no-control rule.

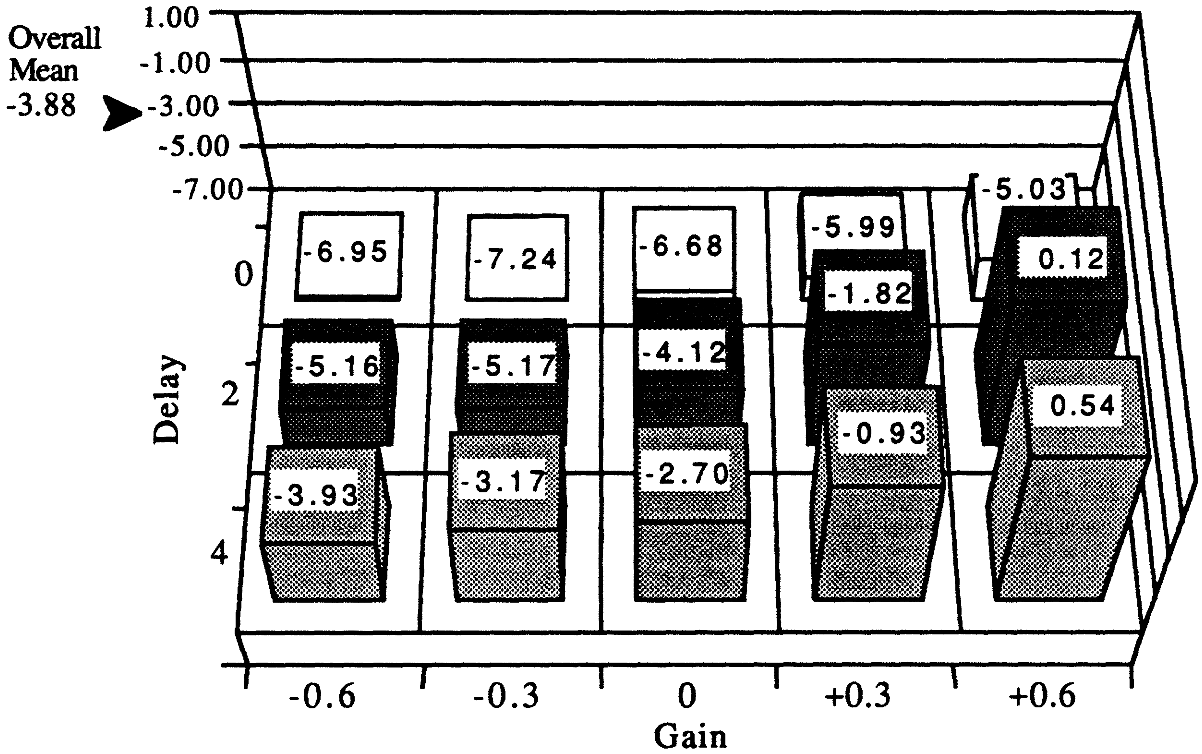

Table 1. ANOVA of $\log_2$(subject costs/no-control costs).

$\overline{R}^2 = .88$, N = 225

Source	df	F	P
Delay	2,182	98.52	p=.000+
Gain	4,182	32.53	p=.000+
Delay x Gain	8,182	2.24	p=.026
Subject	14,182	5.57	p=.000+
Practice/Sequence	14,182	12.88	p=.000+

Table 2 ANOVA of $\log_2$(Subject costs/optimal costs).

$\overline{R}^2 = .42$, N = 225

Source	df	F	P
Delay	2,182	1.02	.364
Gain	4,182	1.22	.303
Delay x Gain	8,182	1.43	.185
Subject	14,182	5.56	.000+
Practice/Sequence	14,182	2.36	.005

Table 3 ANOVA of average time spent on later decisions.

Source	df	F	P
Delay	2,182	1.11	-----
Gain	4,182	0.70	-----
Delay x Gain	8,182	2.67	p<.01
Subject	14,182	15.90	p<.01
Practice	14,182	28.60	p<.01

-- not significant

Table 4 ANOVA of average time spent on first decision in each condition.

Source	df	F	P
Delay	2,182	0.93	-----
Gain	4,182	0.68	-----
Delay x Gain	8,182	1.05	-----
Subject	14,182	3.02	p<.01
Practice	14,182	10.69	p<.01

-- not significant

Table 5. Average adjusted R^2 and incidence of sign reversals for alternative models of subject decisions (see eqns. 9-12).

Model:	M1	M2	M3	M4
Average Adjusted R^2:	0.26	0.51	0.58	0.62
Number of Sign Reversals:	24	59	24	136
Cues Utilized in Each Model:				
I_t	•	•	•	•
$(P_t - S_t)$		•		
$(P_{t+\delta} - S_t)$			•	
P_t				•
P_{t+1}				•
P_{t+2}				•
P_{t+3}				•
P_{t+4}				•
S_t				•

Figure 7 Bootstrap analysis. The total performance shortfall against optimal is partitioned into the effect of inconsistency and the effect of subject's use of the wrong model.

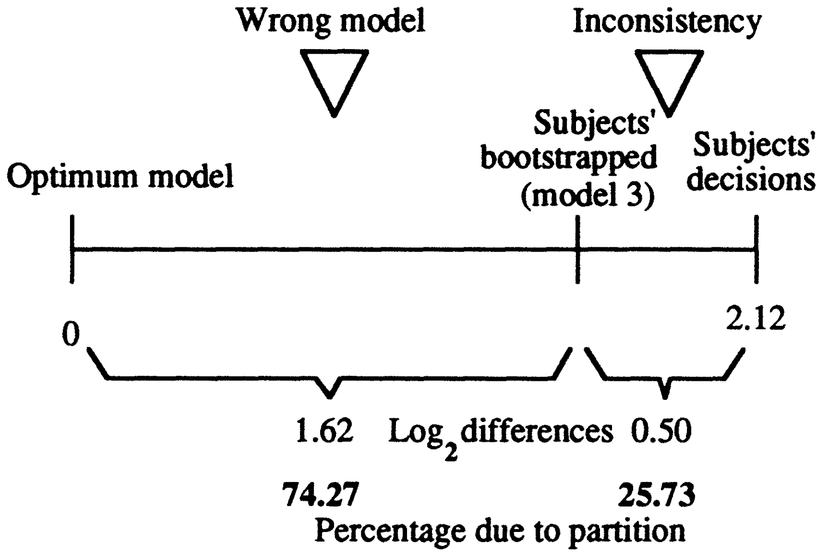

Figure 8. Optimal cue weights for inventory and expected change in inventory compared to the estimated weights, by experimental condition. The weights estimated for each subject are shown as an 'x', and the optimal weights for model 3 are shown as a '+'.

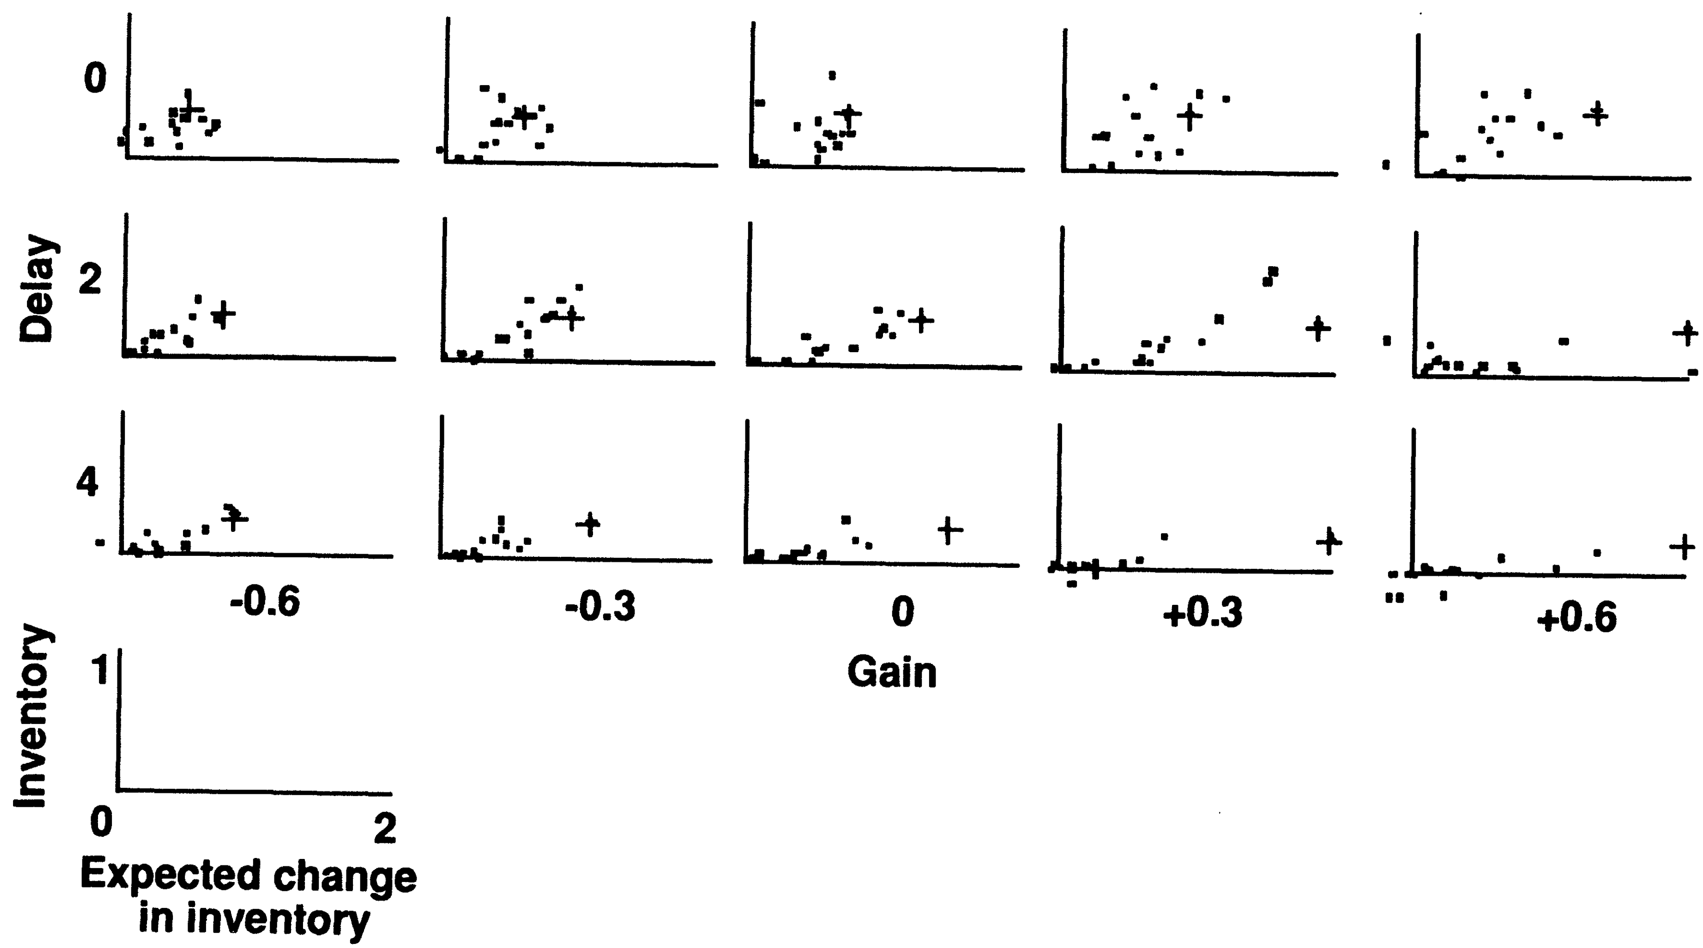

Table 6.

ANOVA for Euclidean distance of subjects' cue weights from optimal for model (see equation 13).

Source	df	F	probability
Delay	2,182	361.25	p=.000+
Gain	4,182	454.07	p=.000+
Delay x Gain	8,182	58.22	p=.000+
Subject	14,182	5.83	p=.000+
Practice	14,182	0.78	------

-- not significant

Figure 9. Dependence of costs on parameters of model 3, by experimental condition. The cost surfaces have been normalized so that each contour line represents a 50% increase in costs compared to optimal. Note the flat optimum in most experimental conditions. Note that for conditions d=2, g = +0.6, and d=4, g = +0.3 and +0.6, the minimum cost is off the scale.