

Building Cognitive Feedback into a Microworld Learning Environment: Results from a Pilot Experiment

Paul A. Langley
London Management Centre
University Of Westminster
35 Marylebone Road
London
NW1 5LS UK

Tel: +44-171-911-5000

Fax: +44-171-911-5059

Internet: langleyp@wmin.ac.uk

12 April 1995

This experimental study examines how performance and learning in a system dynamics microworld environment may be improved through the provision of online cognitive feedback. Subjects are postgraduate management students at the University of Westminster, London. They participate in the experiment over a two week period, as part of a graded assignment. Subjects have to complete a set of six tasks in an Oil Producers microworld, playing the role of the Independents Producers, with a clearly defined performance objective to maximise cumulative net income over a 25 year period. The experimental design includes three different cognitive feedback treatment groups, in addition to a control group which receives no cognitive feedback at all. All groups receive outcome feedback. Treatment groups only have access to the cognitive feedback during the first three trials out of six. Mean subject performance is significantly greater for the treatment groups during the first three trials, but declines to a level comparable with the control group by trial six. Sustainable mean performance improvements are not achieved, but productivity (performance/time taken) *does* improve significantly by the end of experiment.

Introduction

Recent experimental studies conducted within the system dynamics community have shown that dysfunctional behaviour in complex simulated systems may be explained by systematic errors made by the decision makers in failing to account for feedbacks, time delays and nonlinearities (see for example Backken, 1993, Diehl, 1992, Kampmann, 1992). Paich and Sterman (1993) found that performance relative to potential was poor in a simulated complex system involving non-linearities, feedback and time delays. Performance relative to potential was severely degraded when the feedback complexity of the environment was high.

How can we help managers do better, in a microworld learning environment? What tools are needed to help them improve performance? The transfer of learning from

microworlds to the real-world is an important related issue, but is not addressed here. From behavioral decision theory, we note that cognitive feedback -- feedback on "how to" complete a task, rather than just outcome feedback about the results of performance in the task, impacts positively on performance (Balzer et al, 1989). In system dynamics, we frequently describe important cognitive feedback to be the understanding of how systemic structure influences behaviour in complex systems. If we provide cognitive feedback to subjects as they complete microworld tasks, I am interested in how sustainable is the learning that takes place, i.e., when the cognitive feedback is removed what happens to the performance of subjects? Do they suffer *withdrawal* with consequent decline in performance. I ask these specific research questions. When subjects perform complex dynamic tasks in a simulated microworld environment, what is the impact of cognitive feedback on their task performance? Does it make any difference to performance whether the cognitive feedback includes help on decision rules (as investigated extensively in the behavioral decision theory literature), or help on how the task systemic structure impacts on behaviour? After a learning period with cognitive feedback, what happens to subject performance when the cognitive feedback is removed? Does it improve or worsen, relative to subjects who have not received cognitive feedback at all? Do subjects spend longer on the learning activity when cognitive feedback is available? If so, is the extra time justified in terms of performance improvements or learning outcomes?

Experimental Design

The 64 subjects were postgraduate students taking management Masters Programmes in the London Management Centre at the University of Westminster. Subjects undertook the exercise as an assignment as part of a module in Strategic Modelling. The subjects were told that their grade for the assignment would be based on the quality of their logs and write-ups, and not on their performance in the game.

Subjects completed a set of six microworld tasks linked to case study material on the oil industry. The tasks were dynamic in nature and varied in complexity, but were all related to the same Oil Producers system dynamics model (Morecroft and van der Heijden, 1992). The model simulates oil industry behaviour over a 25 year period from 1988 to 2013. It was developed with a team of managers from Group Planning at Shell International Petroleum. It has been published in the academic literature, and has been used in management development programmes with operating company managers from the Royal Dutch/Shell Oil group. It is thus a valid tool for learning about the dynamics of the oil industry.

Subjects performed six tasks, all involving playing the role of the "Independent Producers" in the microworld, over a simulated 25 year period from 1988 to 2013. The tasks were organized as six separate trials, which subjects were allowed to complete over a two-week period. Subjects were allowed to take as long as they wanted over each trial, and were allowed to choose the time elapsed between each trial. Subjects received briefing material which comprised a briefing book (*Oil Producers Microworld -- Independents Game User's Guide*), the Epilogue from *The Prize -- The Epic Quest for Oil, Money and Power* (Yergin, 1991), and recent cuttings about the oil industry from *The Economist* and *The Financial Times*. The User's Guide included background material on oil industry dynamics, a description of important assumptions made in the microworld system dynamics model, a complete set of instructions on how to use the software, a "Getting Started" tutorial for the first game, and a section on "Tips/Tricks" to help subjects remember important points. In addition, all subjects

attended a 1½ hour briefing which covered introductory material about oil industry dynamics, showed 20 minutes of excerpts from *The Prize* (Yergin, 1991) video series, and outlined the game protocols and use of the microworld software. Subjects were required to complete logs before and after trials, which asked them to explain the strategy they intended to use and then to evaluate their results.

Each task involved making yearly oil production decisions, with a performance objective to maximize cumulative net income (cumulative profit) over a simulated 25 year time period. The computer model plays the role of the “Swing” and “Opportunist” oil producer groups, and manages the market structure. The systemic structure of the complex system in which the six tasks are performed is the same, but the tasks differ in terms of exogenous oil demand, and the strategies of the opportunists and independent producers. This ensures that subjects do not repeat exactly the same task twice, and therefore do not benefit from prior knowledge of the system behaviour, i.e., they don’t know exactly what yearly industry demand will be, or the particular strategies of the swing or the opportunist producers.

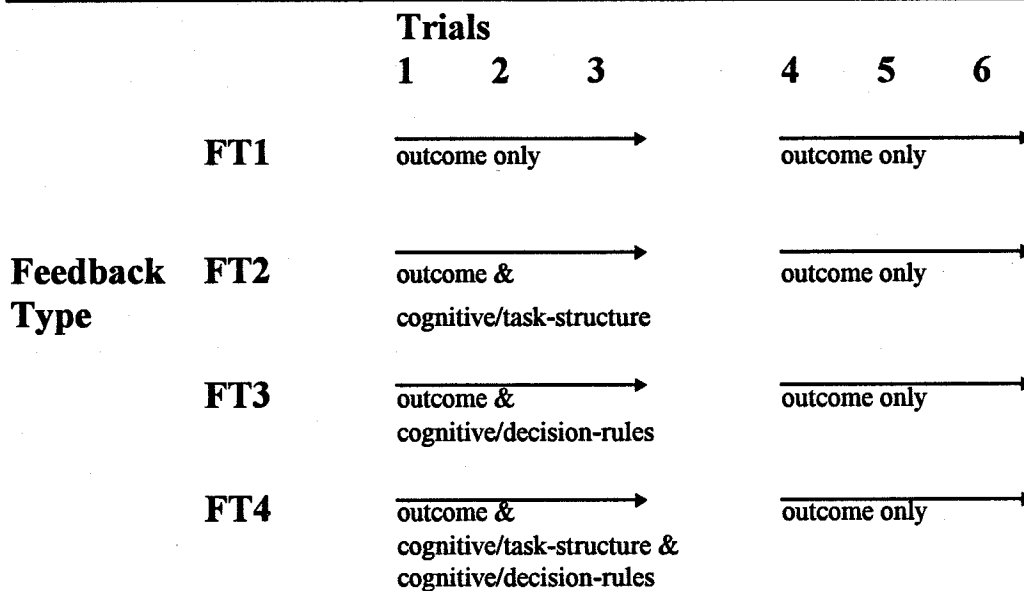


Figure 1 Experimental Design

All subjects performed the same six tasks, in different sequences randomly assigned. The first task (scenario 0) was always the same for all subjects. Subjects were guided through the completion of the first few years of this task by a tutorial in the User’s Guide. The remaining five tasks (scenarios 1 to 5) were completed in a sequence randomly assigned to each subject. There were 120 (i.e., 5 factorial = 120) different possible sequences available. The sequence of tasks was coded onto the subject’s floppy disc, which was needed to run the game software. So, for example, a subject might play the base scenario 0 in trial 1, followed by scenarios 4, 2, 1, 5, 3 for trials 2 to 6 respectively.

The treatment conditions relate to the type of online feedback available to the subjects. Cognitive feedback is divided into two types - *cognitive/task-structure* which links system behaviour to the systemic structure of the task, and *cognitive/decision-rules* which helps the subject formulate decision rules. Subjects receive the feedback treatments during the first three tasks only. Then, they must complete a further three tasks without any cognitive feedback (but

still receiving outcome feedback). Thus the experiment is divided into a *learning phase*, and a *performing phase*, so that we can examine the effect of cognitive feedback on performance during and after the learning phase.

Subjects are randomly allocated to one of the 4 feedback types (see Figure 1):

FT1: outcome only

FT2: outcome & cognitive/task-structure

FT3: outcome & cognitive/decision-rules

FT4: outcome & cognitive/task-structure & cognitive/decision-rules

The experiment is thus a 4 x 6 factorial design, the two factors being Feedback Type (FT1, FT2, FT3, FT4) and Trial (1, 2, 3, 4, 5, 6). Subjects are randomly allocated to one of 4 Feedback Types (FT1, FT2, FT3, FT4), and to one of 120 task sequences for the five tasks (“scenarios”) T1 to T5, performed in Trials 2-6. Note that all subjects perform task T0 in trial 1. The particular sequence in which tasks T1 To T5 is a “nuisance” effect, and is not expected to confound the experimental results. The tasks do vary in difficulty, but have the same protocol. Hence the random assignment of sequences to subjects was chosen in preference to using a Latin Square design (Neter and Wasserman, 1974, p.677). A Latin Square design would have required just five sequences (out of 120 possible sequences) to complete the square with five trials, and a random selection of 5 from 120 is unlikely to be representative. The initial sample size of c.64 subjects provides around 16 subjects for each of the four feedback treatment groups.

Control vs Open Learning

The learning activity is designed to be open-ended, whereby the subjects can take as long as they want to complete the six tasks, within an overall time window of two weeks. I am interested in the time taken to complete the tasks under different feedback treatments, and how this time impacts on performance in the tasks. The microworld software keeps track of these times, as well as the start time for each task so that I can include the time elapsed between tasks in the analysis of variance in subject performance (arguably a time in which subjects *reflect* on the learning activity). One disadvantage of this approach is that I have no control over what the subjects actually do when left to their own devices, or indeed how long they spend doing it. Individual subject effort may vary, introducing additional between-subject variance.

Anything that contributes to the learning activity is encouraged, but a number of actions could be undertaken which would bias the results. Possibilities include performing extra trials, colluding with expert friends, and tampering with the stored data on the discs. The discs were security coded with checksums (on the subject number, sequence number, and cognitive feedback treatment type) to try and prevent most types of data disc tampering. In the briefing sessions, subjects were encouraged not to tamper with the discs. They were also told that the task scenarios were all different (strictly speaking not correct), and no advantage could be gained from working with other subjects. Additionally, they were also told that each time the gaming software was executed, the subject number, date and time were logged in network files. Hence the microworld software would be able to detect if a subject aborted a run before the end of the 25 year simulation period.

Design of the Cognitive Feedback

The cognitive feedback is divided into two types - *cognitive/task-structure* which links the system behaviour to the systemic structure of the task (cause-effect relationships between variables), and *cognitive/decision-rules* which helps the subject formulate decision rules through the provision and explanation of a benchmark rule.

The approach taken is to present the user with a simplified feedback loop structure diagram, indicating the direction of change between variables in the loops using "O" and "S" symbols (see Senge, 1990 for examples). The user can click on any variable in the diagram, and receive advice on the system dynamics that will tend to change the magnitude of the variable. The philosophy is to give a "notion" of cause and effect relationships, without revealing algebraic relationships. For example, if users click on **Production Shortfall**, they will see the following pop-up window on the right of the screen (Figure 2).

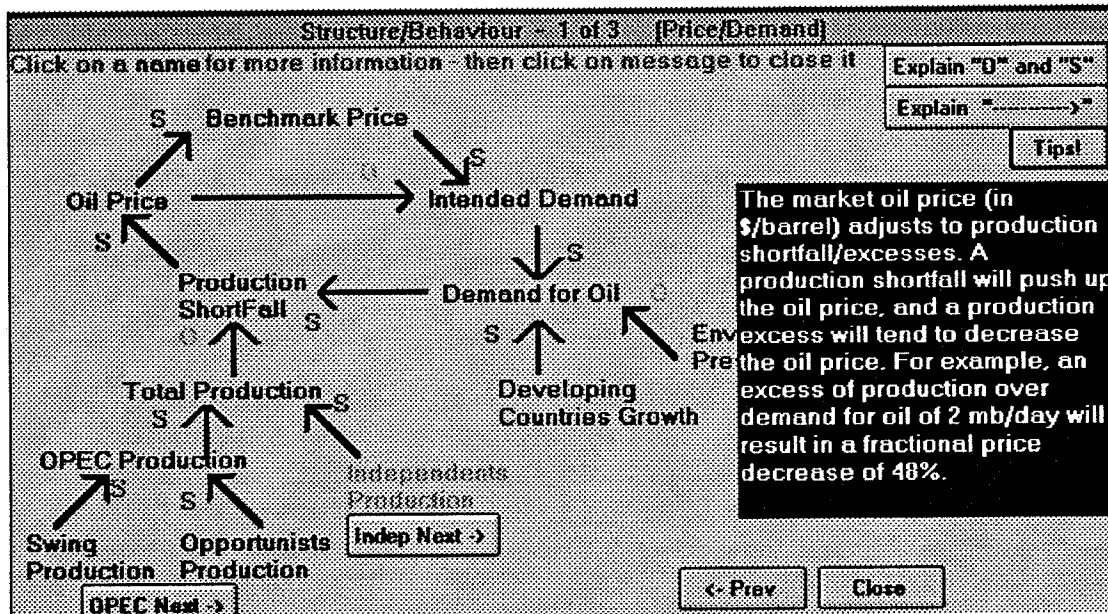


Figure 2 Pop-up window (screen right) for the variable "Production Shortfall"

There are also some general "Tips" available relating to the dynamics of this part of the system, presented in the form of "questions and answers" on a pull-down menu. By clicking the Tips button on the right of the screen, the user will then see a pull-down menu of questions (Figure 3).

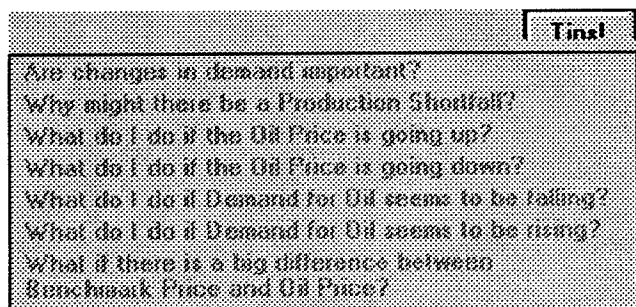


Figure 3 Menu of Questions Available as Tips

On selecting a particular question by clicking on it, a pop-up help window appears. Again, the philosophy is not to give a direct answer to the question, but to suggest to users how they might find the answer. For example, the help window presented to users in response to the question **Why might there be a Production Shortfall?** (Figure 4) suggests that the OPEC Meeting report (outcome feedback) may be checked to help answer that question.

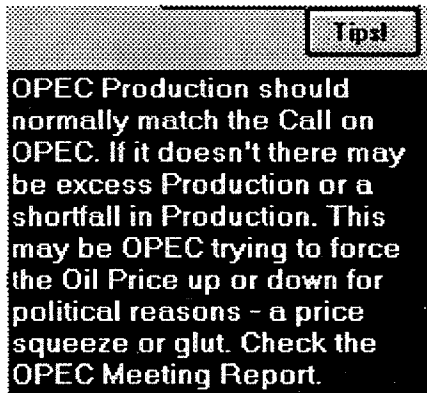


Figure 4 Response to Question about Production Shortfalls

Moving now to the cognitive feedback/decision rules screens. The second screen (Figure 5) shows the non-linear relationship between the **Viable Fractional Increase in Capacity** and the **Profitability Ratio** (which is essentially the ratio of **Expected Oil Price/Development Cost per Barrel**). It also suggests that the user bear in mind the depletion rate of existing capacity, when considering how much new capacity to approve.




Figure 5 Cognitive Feedback/Decision Rules Screen 2 of 3

The final screen three (Figure 6) allows the user to make a judgment about investment optimism. The recommended capacity approval is calculated for a variety of different Expected Oil Prices, and Hurdle Rates. The Hurdle Rate is the analogy of the Independent

Producers' weighted average cost of capital. An optimistic investor may take a lower hurdle rate than the default 0.15, and maybe also be forming expectations of a higher oil price than the project appraisal team. Conversely, a pessimistic investor may choose a higher hurdle rate than the default 0.15, and may be forming expectations of a lower oil price than the project appraisal team. The final decision is left to the user. The user is expected to interpolate values between the indicated values specified on the screen.

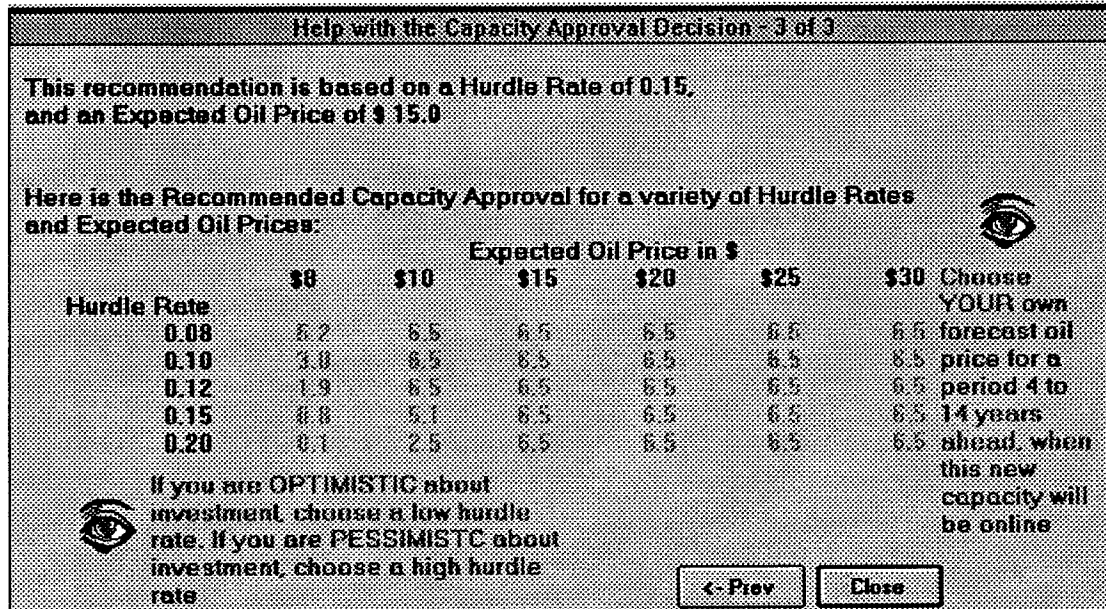


Figure 6 Cognitive Feedback/Decision Rules (FT3) Screen 3 of 3

Results

Figure 7 shows how the mean subject Performance Relative to Benchmark (PRB) -- the mean ratio of subject cumulative profit to the benchmark profit -- varies over the six trials and four feedback treatment conditions. The mean subject performance in the tasks is below the benchmark (of 1.0) in all trials, and across all four feedback treatments. The control group (FT1), who received outcome feedback only, performs worse in trials 1-3 than the other three treatment groups ($F(3,383) = -3.28, p < 0.001$). However, the control group FT1 improves steadily over time, and by trial 6 the mean performance is close to groups FT2, FT3, FT4. Group FT3 (cognitive feedback/decision rules) does the best in trial 1 ($F(15,383) = 3.21, p=0.001$) and 2. The initial higher performance of the treatment groups FT2, FT3, FT4 seems to level off/decline slightly in trials 4,5,6 when the cognitive feedback is no longer available. As we might expect, the treatment group with cognitive feedback/task structure assistance (FT2) seem to maintain their level of performance attained in trials 1-3, whereas the group with cognitive feedback/decision rules assistance (FT3) declines in performance. Overall, the outcome feedback only control group (FT1) seem to do as well after 6 trials as the other treatment groups. The treatment groups' (FT2, FT3, FT4) early lead is lost in the last three trials 4,5,6. If they had continued to receive the cognitive feedback during trials 4,5,6 would their performance have continued to increase? The provision of cognitive feedback accelerated the *surface* learning about the tasks, but seemed to fail to achieve the *deep* learning necessary to achieve sustainable performance improvements.




Figure 7 Graph of Mean Subject Performance Relative to Benchmark (PRB) in Trials 1-6 under 4 Feedback Treatments FT1-FT4

Percentage of Subjects who Beat the Benchmark

Figure 8 shows how the percentage of subjects who “beat” the benchmark varies across six trials and four feedback treatments. Despite the fact that the *mean* subject performance in all six trials and in all four feedback treatments was below the benchmark, an impressive number of subjects manage to beat the benchmark, as they learn to improve performance. FT2 (cognitive feedback/task structure) does the best -- in trials 3 and 4, 56% of the subjects perform better than the benchmark. But, by trial 6 the percentage is reduced to 35%. Certainly, it appears that the higher performers do very well under feedback treatment FT2, maintain this high performance for one further trial, then collapse. As before with the mean performances, by trial 6 there are only small differences between the control and treatment groups.

Time Spent on the Task and on the Cognitive Feedback Screens

In Trial 1, the time spent by subjects on completing one task (i.e., making 25 yearly decisions for capacity approval) varies from 53 seconds, to 16 hours 10 minutes! The mean time for trial 1 is 80 minutes (FT1=95.4 mins, FT2=58.3 mins, FT3=123 mins, FT4=43.5 mins). Times decrease significantly over the six trials. By trial 6, the mean time is 18.8 mins (FT1=20.1 mins, FT2=25.3 mins, FT3=18.6 mins, FT4=11.7 mins).

In Trial 1, the time spent by subjects looking at cognitive feedback screens for subjects in feedback treatment groups FT2, FT3, FT4 varies from 0 minutes (i.e., chose not to make use of the cognitive feedback) to 61 minutes, with a mean of 7.2 minutes. By trial 3, this mean has reduced to 1.4 minutes. The longest times in all three trials are for group FT3 (cognitive feedback/decision rules). Overall, the time spent studying cognitive feedback screens is surprisingly low.




Figure 8 Graph of Percentage of Subjects Exceeding the Performance Benchmark (PRB), in Trials 1-6, under Feedback Treatments FT1-FT4




Figure 9 Graph of Performance Relative to Benchmark (PRB) divided by Time Spent on Completing Task (T_s) in Trials 1-6, under 4 Feedback Treatments FT1-FT4

Subject Productivity

One measure of subject productivity is the ratio of (performance relative to benchmark / time spent on task). Figure 9 shows how this productivity ratio varies across trials for the four feedback treatment groups. Group FT4 (decision rules & task structure) does very well indeed, outperforming the other three groups in all trials apart from trial 3 ($F(3,383) = 2.8$, $p = 0.012$). Productivity continues to improve even after the cognitive feedback is no longer available in trials 4,5,6. The implication here is that although the mean and upper-quartile subject performance of the cognitive feedback groups is not significantly better by trial 6, than the outcome feedback only group (FT1), the time taken to achieve similar performance is less.

Discussion

At the time of writing, work is very much still in progress with further sets of subjects. Written protocols (assignment logs) are being analyzed for insight into which information cues subjects found important, and whether these cues changed over time (as subjects learned). In order to explore further the issue of whether performance may be further improved if the cognitive feedback is available for later trials, I am providing a treatment group with cognitive feedback for all six trials.

References

- Bakken, B.E. 1993. *Learning and Transfer of Understanding in Dynamic Decision Environments*, PhD dissertation, Sloan School of Management, MIT, Cambridge 02142.
- Balzer, W.K., M.E. Doherty and R. O'Connor Jr. 1989. Effects of cognitive feedback on performance. *Psychological Bulletin*, 106, 410-433.
- Brehmer, B. 1990. Strategies in Real-Time, Dynamic Decision Making, in R. Hogarth (Ed.), *Insights in Decision Making: A Tribute to Hillel J. Einhorn*, Chicago, IL: University of Chicago Press.
- Diehl, Ernst W. 1992. *Effects of Feedback Structure on Dynamic Decision Making*, PhD dissertation, Sloan School of Management, MIT, Cambridge MA 02142.
- Kampmann, C.P.E. 1992. *Feedback Complexity and Market Adjustment: An Experimental Approach*, PhD dissertation, Sloan School of Management, MIT, Cambridge MA 02142.
- Morecroft, J.D.W. and K.A.J.M. van der Heijden. 1992. Modelling the Oil Producers - Capturing Oil Industry Knowledge in a Behavioural Simulation Model. *European Journal of Operational Research*, 59 (1), 102-122.
- Neter, J. and W. Wasserman 1974. *Applied Linear Statistical Models*, Homewood, Ill: Irwin.
- Paich, M. and J. Sterman. 1993. Boom, Bust, and Failures to Learn in Experimental Markets, *Management Science*, 39, 12, 1439-1458.
- Sengupta, K. and T. Abdel-Hamid 1993. Alternative Conceptions of Feedback in Dynamic Decision Environments: An Experimental Investigation, *Management Science*, 39, 4, 411-428.
- Yergin, Daniel 1991. *The Prize -- The Epic Quest for Oil, Money and Power*, London: Simon and Schuster.