Exploration of Outcome Feedback for Dynamic Decision Making

Naiyi Hsiao

Department of Public Administration and Policy
Nelson A. Rockefeller College of Public Affairs and Policy
University at Albany, State University of New York

Center for Policy Research
Milne 300
135 Western Avenue
Albany, New York 12222
USA

+1 518 442 3852
nh7365@csc.albany.edu
http://ALPHA1.albany.edu/~nh7365/prospectus.htm

ABSTRACT

The research proposes alternative designs of outcome feedback incorporating the two components of outcome feedback that have been missed in the literature of dynamic decision making. It has been theoretically accounted and hypothesized that outcome feedback with the two components –benchmark outcome and benchmark decisions – will help decision makers manage dynamic decision tasks more effectively.

Three treatment groups have been set up to empirically test the hypotheses in a gaming experiment built on the Beer Game. The first is a group of subjects who receive knowledge of results – the actual results of their own decisions – alone. The subjects in the second treatment receive benchmark outcome, perceived as the best competitor's outcome, aside from knowledge of results. The subjects in the last group have access to full-featured outcome feedback – knowledge of results, benchmark outcome, and benchmark decisions – perceived as the best competitor's decisions – while they are managing the game.

1. INTRODUCTION

Three types of information feedback in DDM environments have been proposed in the literature: outcome feedback, cognitive feedback, and feedforward – as reviewed in the next chapter. Based on the studies, outcome feedback generally refers to the information feedback that displays the results of the decisions being made previously. Cognitive feedback covers the broader sense of task related knowledge such as the structure of the task, numerical relationship among variables, and feedback loops understanding. Feedforward, or decision heuristics, is a simple set of rules that may help decision makers produce effective decisions. These three types of information content share the purpose of helping decision makers deal with dynamic decision tasks more effectively. The DDM studies have been experimenting with various conceptualizations and designs of these three types of information feedback and empirically examine their effectiveness as decision aids.

Some terminology used throughout the research should be clarified here. First of all, a decision variable stands for a set of variables that decision makers can manipulate to achieve some pre-set outcome measures – the evaluative criteria for human decision-making behavior. In the context of dynamic decision making, this set of decision variables has to be determined over some pre-defined time frame. Secondly, the distinction between outcome feedback and other available information involved in a dynamic decision task is important. Outcome feedback is simply the results, in terms of various outcome measures, of the decisions that have been made previously. As will be presented soon, the current study attempts to add some elements to this conceptualization. Available information includes all task variables other than decision variables and outcome measures. Studying these items of accessible information may help decision-making. "Cues" will be utilized throughout the research to stand for these various items of available information.

2. LITERATURE REVIEW AND REFLECTIONS

The conceptual definition of outcome feedback is then investigated in the literature of judgment and decision making – the Brunswikean lens model provides an explanatory view of outcome feedback. The comparison of how outcome feedback is treated in the different literature motivates the thinking about how decision makers apply their mental models in handling dynamic decision tasks. This thinking and subsequent hypotheses development turn to a set of working hypotheses to examine two alternative designs of outcome feedback.

2.1. A Literature Review of Dynamic Decision Making

A comprehensive literature review of dynamic decision making has been done by the author of the current research (Hsiao, 1998). Five groups of dependent variables (evaluative criteria for dynamic decision-making behavior) and three groups of independent variables (predictors for dynamic decision-making behaviors), as well as a tentative framework to illustrate these predictors, have been identified based on the 33 reviewed studies that conduct gaming experiments on various tasks. The brief summary below only provides a background for the ongoing exploration of outcome feedback. Refer to the original work (Hsiao, 1998) for more findings for the literature review.

Laboratory experiments have employed a wide range of tasks in the literature of dynamic decision making. They include the sugar production task (Berry and

Broadbent, 1984), which involves only a few variables and a single decision, and the welfare system (Maxwell, 1995), which requires six decisions being made simultaneously and more than a hundred variables. These microworlds (or flight simulators, learning laboratories, learning environments) share the following general characteristics (Brehmer and Dorner, 1993): (1) complex – decision makers are required to pay attention to many things, such as goals, side effects, and choices among different actions, (2) dynamic – decision makers have to perceive task systems changing over time both as a consequence of their decisions and autonomously, (3) opaque – to perform and learn well decision makers must be able to infer task structures and develop their decision strategies.

In sum, there have been three types of information contents that the previous studies have attempted to provide decision makers as decision aids: feedforward, cognitive feedback, and outcome feedback. Various forms of feedforward (decision heuristics) and cognitive feedback have been experimented. Although their effectiveness is mixed, both are expected to enhance decision makers' understanding of the task knowledge (declarative knowledge) and decision rules development (procedural knowledge), which would hopefully lead to better task performance.

The current research, based on the literature review, starts with an interesting observation that outcome feedback has been taken for granted by the literature without formal tests for its effectiveness, except Sengupta and Abdel-Hamid (1993) concluding that providing additional information about decision rules and task structure improves task performance compared with providing outcome feedback – the results of the previous decisions – alone. The lack of more formal examination of outcome feedback largely results from the fact that outcome feedback – the outcome of the decisions being made – is one of the defining characteristics of dynamic decision making. In other words, outcome feedback has always been available in all DDM studies. Moreover, some studies (e.g., Sterman, 1989a, 1989b) have attempted to provide complete information, including outcome feedback and status reports for all task variables in both numerical and graphical formats, and found that decision makers still suffered from ill performance. In sum, the literature seems implicitly to conclude that outcome feedback alone is not effective as a decision aid.

Nevertheless, the following two threads of thinking motivates the current study to explore outcome feedback. The first reviews the concept of outcome feedback from the literature of psychological judgment and decision making – particularly the literature of social judgment theory (Cooksey, 1996). The second stems from the reflections on the information decision that makers may have in dealing with dynamic decision tasks in the real world.

2.2. Outcome Feedback – an "Alternative" View from another Literature

The literature of judgment analysis (Cooksey, 1996) provides a conventional and clear view for the concept of outcome feedback. As Figure 1 below, the lens model framework has attempted to describe the elements of a typical judgment task. At first, experts make judgment (Y in Figure 1) on some task, such as weather forecasting, based on some information items – the cues (Xi). Then there will be the correct answer or standard for the judgment being made, termed criterion (O). There always involves errors on both sides – human judgment and task environments (Eo and Ey). In this context, the outcome feedback of the human judgment (Y) is defined by the criterion (O).

The correlation between the human judgment and outcome feedback is described by a measure called judgment achievement (r_a) . Further, as the human judgment is based on the cues, the correlation between the cues and the judgment (Rs) stands for how consistently a human judge can apply his/her judgment policy. Similarly, the criterion is also related to those cues with the correlation (Re), which stands for how well the criterion can be predicted by a model based on the cues and therefore termed task predictability.

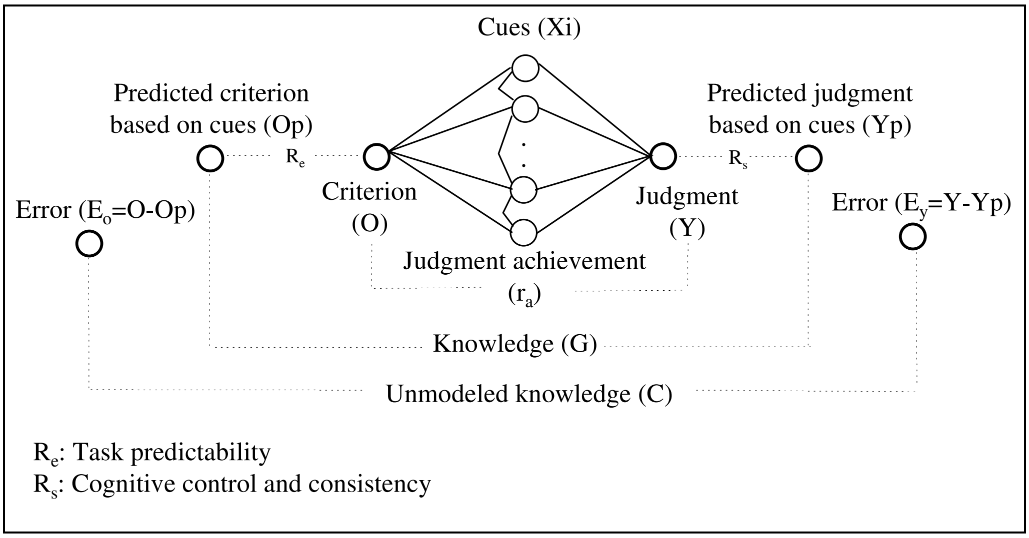

Figure 1: The lens model framework

Note the critical components of outcome feedback according to the lens model framework. The criterion is the standard for the human judgment being made. If a meteorologist predicts whether or not it will rain tomorrow, the criterion would be the actual weather condition tomorrow, which will be known the next day. There are two components that are usually mixed together underlying the concept of outcome feedback. First, the criterion, actual weather tomorrow in this case, stands for the ***correct outcome*** of the weather forecast. Second, this criterion meanwhile suggests that this same prediction should have been made by the meteorologist, which leads to the sense of ***correct judgment***. This level of conceptual decomposition of outcome feedback is critical. However, it has been unfortunately unattended and even ignored by the literature of dynamic decision making, most of which equate the term with the root in the decision-making literature– outcome feedback – but fail to explore the two underlying components specified above.

As the literature review summarized above, outcome feedback in the literature of dynamic decision making is measured by the actual outcome of the decisions being made. While the measure and conceptualization appear consistent with the meaning of the word "outcome," they miss the very two components underlying the original concept of outcome feedback analyzed above. In the first place, the "actual outcome of previous decisions" contains no indication of either "standard" or "correct" outcome; rather, it is just the results of the human decisions. Secondly, outcome feedback employed in most DDM studies loses its suggestion of correct judgment – "***correct decisions***" in the context of dynamic decision tasks.

The somewhat distorted use of outcome feedback in dynamic decision making environments, nevertheless, has its legitimate explanation. Comparing the context of a typical judgment task (as depicted in the lens model framework) and that of a dynamic decision task reveals the fundamental distinction. In a typical judgment task, the task system (the atmospheric phenomenon in the preceding example) containing the variable being judged (the weather tomorrow) by human experts is not affected by the expert judgment. To illustrate, whatever a meteorologist forecasts will not change the actual weather tomorrow. This contrasts with one of the defining characteristics of a dynamic decision task – the task system changes both autonomously and because of the decisions being made. Accordingly, it seems natural to conceptualize outcome feedback as the actual outcome determined by the decisions for dynamic tasks.

Nevertheless, the missing components – correct outcome and correct decisions – can still make sense in dynamic decision environments provided that a broader view is adopted. If correct outcome is conceived as the results from correct decisions and "correct" is conceived as "optimal," then correct outcome in the context of dynamic decision making stands for the best outcome that can be achieved by optimal decisions. While the rephrase of outcome feedback appear promising, it entails the existence of optimal decisions for dynamic decision tasks. It has long been puzzling as to the conceptual definitions of optimal solutions for a dynamic decision task (Andersen and Rohrbaugh, 1992). Further, even though the conceptual definitions can be settled, most dynamic complex problems are extremely difficult, if not impossible, to derive optimal decisions analytically. The best the research work has been attempting is to derive some decision rules – benchmarks – that, given the parameters properly assigned, can produce better outcome compared with human decision makers in most cases.

The current research builds on the foregoing conceptual review and argues for the incorporation of the missing components of outcome feedback. The research questions of the central interest are: Will the outcome feedback, with the two missing components, still work poorly as decision aids as most DDM studies conclude in the literature? Will the full-featured outcome feedback be of any help for resolving dynamic decision tasks? If so, what would be the theoretical account and how could empirical evidences be collected? These research questions will make even more significance when coupled with the following reflections on dynamic decision making environments in the real world.

2.3. Reflections on the Real-World Dynamic Decision Making

Observing how decision makers handle dynamic decision tasks in the real world further encourages the use of the full-featured outcome feedback proposed by the current research. Consider what decision makers can possibly have in hands when they face a dynamic decision problem: (a) understanding – task knowledge and decision heuristics – about the decision problem and increased expertise and experience, (b) actual results based on their previous decisions - knowledge of results hereafter, i.e., outcome feedback defined by most studies of dynamic decision making as indicated above, (c) various information items (cues) that help decision-making, usually with the limited accessibility, and (d) with luck and substantial delay, the decisions and outcome of their competitors in dealing with the same problem.

If the observation captures some reality, the real world hardly provides decision makers with cognitive feedback and decision heuristics which have generally been more effective compared with outcome feedback evidenced by most studies. Rather, decision makers acquire task knowledge and develop decision heuristics to accumulate the stock of expertise as in (a) above based on what they have and apply effectively on knowledge of results (b), cues with limited availability (c), and competitors' decisions and outcome with delay (d).

This reflection on the real world, aside from the conceptual exploration of outcome feedback, further stimulates the current research: How should outcome feedback be designed in order to help decision makers more effectively acquire task understanding and/or develop useful decision heuristics, and hence improve task performance?

3. HYPOTHESES DEVELOPMENT

The literature review summarized in the foregoing discussion has revealed that task complexity has been an important factor lurking behind the effect of any decision aids as indicated above. To illustrate, any design of decision aids may not be working as expected if decision makers face a dynamic decision task with extreme difficulty. This is particularly likely for some studies based on a simulation model containing hundreds of variables, requiring multiple decisions being made, and evaluating task performance measured by multiple mutually-conflicting outcome measures, such as Jansson (1995). In the other extreme, the effect of decision aids may also be concealed if decision makers can easily resolve an easy task; that is, no decision aid is necessary because the task is so simple. To avoid the possible lurking effect of task complexity, the research proposes to examine the foregoing designs of the outcome feedback in two levels of difficulty for the same task. According to all exposition above, the following Table 1 provides an experimental design to test the effectiveness of alternative designs of outcome feedback.

Three treatment groups represent various designs of outcome feedback: knowledge of results as in the current literature (Treatment Group A), knowledge of results with the addition of benchmark outcome (Treatment Group B), and knowledge of results with the addition of benchmark outcome and decision (Treatment Group C). The treatments will be tested under two levels of tasks with manipulated difficulty detailed in the next chapter. Although no specific hypotheses have been proposed about the effect of levels of task complexity, this dimension avoids the possible lurking effect on the designs of outcome feedback as argued above.

Table 1: Design on Treatment Groups and Task Difficulty

	Treatment A: Knowledge of Results (KOR)	Treatment B: KOR plus benchmark outcome	Treatment C: KOR plus benchmark outcome and decision
Task Level 1 (The easier task)	A1	B1	C1
Task Level 2	A2	B2	C2

Also based on the evaluative criteria for human decision-making behavior (Table A3 in Appendix A), the effectiveness of the three designs of outcome feedback will be primarily evaluated on task performance (A1 in Table A3), task knowledge acquisition (declarative knowledge, B1 in Table A3), and decision heuristics development (procedural knowledge, B2 in Table A3). A list of hypotheses based on the foregoing theoretical exploration on outcome feedback as decision aids is provided here.

(A1) Subjects in the treatment group B (knowledge of results plus benchmark outcome) outperform those in the group A (with knowledge of results alone) in terms of task performance.

(A2) Subjects in the treatment group B outperform those in the group A in terms of knowledge acquisition.

(A3) Subjects in the treatment group B outperform those in the group A in terms of heuristics development.

(B1) Subjects in the treatment group C (knowledge of results plus benchmark outcome and decision) outperform those in the group A in terms of task performance.

(B2) Subjects in the treatment group C outperform those in the group A in terms of knowledge acquisition.

(B3) Subjects in the treatment group C outperform those in the group A in terms of heuristics development.

There are another evaluative criteria for the effectiveness of the various outcome feedback. The following two measures are those the current research chooses to look at: amounts of decision time (C1 in Table A3) and amounts of decision use for specific information items (C2 in Table A3). The following hypotheses are proposed based on the theoretical account that the subjects with the benchmark information will have more information to explore as reasoned above

(A4) Subjects in the treatment group B spend more decision time than those in the group A.

(A5) Subjects in the treatment group B have more access to specific information items than those in the group A.

(B4) Subjects in the treatment group C spend more decision time than those in the group A.

(B5) Subjects in the treatment group C have more access to specific information items than those in the group A.

4. A PRODUCTION- DISTRIBUTION TASK - THE BEER GAME

The generic production-distribution task dates back to the early development of the field of system dynamics modeling and has been refined for the current edition of the Beer Distribution Game. See Sterman (1992) for an annotated bibliography and the game development. Figure B1 in Appendix shows the original game board. In the board game, four players are involved in the production-distribution system – factory, distributor, wholesaler, and retailer. Each of the players receives the beer ordering from its downstream player and place ordering accordingly from its upstream player. The players, through a series of beer ordering from the upstream, attempt to manage their beer inventories. A high level of inventory incurs inventory cost, and an out-of-stock condition incurs backlog cost as well. The objective for each player, as well as a team with the four roles, is to minimize total costs over a pre-set period of time, 25 weeks in the present study. In addition, no communication between the four players is allowed. That is, each player only has access to its own part of information.

Note that the board-game edition of the Beer Game meets the three defining characteristics of dynamic decision making mentioned above. It also adds another factor of complexity with the interaction between players – decisions made by any player change the environment and hence affect the decision of the other players.

Sterman (1989b) has employed this edition of the Beer Game to illustrate the decision makers' misperceptions of feedback in the dynamic decision experiment.

The current research, based on the board-game edition, revises the task in the following manners. First, the production-distribution system underlying the Beer Game is transformed into a computer simulation model which contains the identical set of equations as in Sterman (1989b). This implies that players have to deal with the computer interface, rather than the game board and chips, in order to manage the task. The design of the computerized gaming interface becomes an important issue to avoid an unfortunate condition that players fail to manage the task because they fail to interact with the computerized gaming interface where all information is displayed. It deserves emphasis again that the information display is absolute critical and can hardly be separated from the information contents. As the central purpose of the current work is to test the various contents of outcome feedback, the information display has to be controlled to minimize its impact on the experimental results. This point will be elaborated later in the section of the interface design.

Second, corresponding to the requirement of two levels of task difficulty, this study utilizes the same Beer Game with two levels of complexity by manipulating (a) the degree of random variation of the environments – the weekly sales to the beer consumers in the Beer Game – in that higher random variation suggests a more difficult task as in Mackinnon and Wearing (1980), (b) the time delay or decision effectiveness – the production and shipment delays in the Beer Game – in that longer delays lead to a more difficult task as evidenced by many studies such as Kleinmuntz (1985) and Diehl and Sterman (1995), and (c) the system equations that change the strength of positive feedback loops (the pre-defined decision heuristics of players in the Beer Game) in that stronger positive gains cause more drastic oscillation patterns as in Paich and Sterman (1993).

Particularly, all subjects in the gaming experiment will play the beer retailers. The single decision required is to place the weekly beer ordering from the upstream wholesaler. The simulation model will produce the decisions for the other three sectors: wholesaler, distributor, and factory, and the two levels of task complexity is partly achieved by manipulating the decision heuristics that these three sectors apply. Task performance is represented by the beer inventory and the costs. A positive amount of inventory causes the inventory cost; a negative amount of inventory incurs the out-of-stock cost or backlog cost. The sum of the inventory cost and backlog cost is termed as the total cost. The goal of the retailer is to maintain the weekly total cost as low as possible so that the total cost over the 25-week gaming period may be minimized. In this case, a better decision / outcome in a short term will lead to better outcome in the long term. This reduces the goal complexity although it might be unusual in a typical complex dynamic task.

The single decision and the coherent short term vs. long term performance measures primarily account for the choice of the Beer Game for the current research. The performance measures – weekly (short-term) and total (long-term) inventory / cost – are apparently understandable and therefore leave few chances for decision makers to wrongly evaluate the outcome of their decisions. Therefore, based on the argument in the hypotheses development above, the discrepancies between subject groups in terms of task performance, task knowledge and heuristics formation can be attributed to the various designs of outcome feedback rather than to the subjects failure to evaluate the outcome of their decisions.

In addition to their decisions (the weekly beer ordering) and the outcome (the inventory and cost), subjects have also access to two items of relevant information.

The first cue is the weekly sales to the beer consumers. The second cue is the arrivals of beers from the wholesaler which respond to previous beer orders. As retailers, subjects rely on both cues to estimate the next beer order. Note that only these two cues are available for subjects playing retailers although there are other relevant cues in the task system they do not have access to. The limited access to relevant cues is consistent with the rule of no communication between sectors in the original board game. In addition, it also avoids the information-search failure mentioned above – decision makers may look at irrelevant cues to make decisions. Similar to the design of the decision and performance measures above, this increases the explanatory power of the examined designs of outcome feedback for the experimental results. To sum up, discrepancies of subject groups on task performance, task knowledge and heuristics development can be more intimately related to the differential designs of outcome feedback, through subjects capability to establish the decision-outcome relation, cue-decision relation, cue-outcome relation, and the mental model of the relations.

The heuristics for benchmark decisions is developed based on the analysis of the generic production-distribution task in Sterman (1989b). Basically the benchmark decision heuristics utilizes the pipeline information that takes into consideration of the delayed effect on production and distribution and hence avoids over-stocking and reduce oscillation of the beer inventory. The benchmark outperforms human decision makers based on the empirical evidences in most, if not all, circumstances. Thus it should served well as the market competitor that may stimulate subjects to pursue and experiment their thinking about the difference between their decision / outcome and the benchmark.

5. EXPERIMENTAL PROCEDURES

Table 2: A Summary of the Experimental Procedures

Step	**Brief Description (details below)**	**Time Estimated (minutes)**
1	- Check experiment materials - Program installation	60
Subtotal for Step 1		**60**
2	- Announce the schedule - Pre-game training - Subjects practice the interface	40
3	- The formal game starts - At least three trials (each with 25 decisions) without formal break	150
Break		10
4	- Post-game debriefing	40
Subtotal for Step 2-4		**240**
5	- The session dismissed and subjects get paid before they leave	10
6	- Collect questionnaire and data	50
Total for Step 1-6		**360**

Here is a step-by-step list ofactivities in each session of the gaming experiment. The following Table 2 summarizes the sequence of activities in the session.

Step 1: Make sure all materials required in the experiment are prepared. This includes the hardware and gaming program tests, notepads, pens, signed consent forms, and envelops with the subjects remuneration. This would take the researcher around an hour before subjects start gathering.

Step 2: The pre-game training session starts with the agenda for the whole gaming session, including the schedule for all activities below. Those subjects who have not signed the consent form will be asked to sign if they agree to voluntarily participate in the gaming experiment. Then the researcher will guide subjects through the paper documents about the Beer Game, the task information, and the sample screens as shown in Appendix C. They will also be informed that these pre-game training materials are always accessible anytime they play the game. The subjects can then start practicing the game. The researcher will tutor the first 3 decisions to make them familiar with the operation of the program, the interpretations for graphs and tables, and the decision being made. They then will be given around 10 minutes to practice the game. Totally it would take around 40 minutes for the whole training session. In addition, subjects will be encouraged to take notes while they are playing. They will be told that the notes will be very helpful for them to answer the post-game debriefing questions.

Step 3: The formal game. The researcher will announce the formal start of the gaming task. The subjects will also be told that 40 minutes should be sufficient for each trial with 25 decisions to be made – given that each decision will take around 60 to 90 seconds and it will usually take less when the subjects get familiar with the gaming interface. They are asked to complete at least three trials of the game in the next 150 minutes and encouraged to perform as many trials as they can – as long as they will not make rush decisions without exploring the task. There will be no formal break in this step and no communication about the gaming task between subjects is allowed. The researcher will remind them when there are 20 minutes left. When time is up, all subjects will be asked to stop the game immediately. Then there will be a formal 10-minute break.

Step 4: Post-game debriefing starts after the preceding 10-minute break. The purpose of the debriefing is to extract the thinking that runs through the subjects mental activities. All written questions in the debriefing, introduced in the next section, are designed to measure whether the subjects acquire task related knowledge and develop decision heuristics. Although no communication between subjects is allowed, they are encouraged to answer the debriefing questions referring to their notes taken when they were playing the game. The subjects will be told that they have 40 minutes to finish all questions and an announcement of time will be made by the experimenter. See details in the next section.

Step 5: The session will be dismissed and the subjects can get their economic reward $32 for each individual before the leave. As the payment has been in envelops already prepared before the session (Step 1 above), it should take less than 10 minutes for 10-12 subjects each session.

Step 6: Materials collection. The researcher will now have to collect all written questionnaires, notes, and data from the individual computers. The system will be designed in a way that the data will be collected in the files with specific identification. This will make data easier to be collected and sorted out.

Note that the subjects take part in the gaming session from Step 2 to Step 4, totally taking them around 4 hours. However, it will take extra hour for the preparation (Step 1) and collection (Step 5 and 6) respectively for the experimenter. This will require a 6-hour reservation of the electronic classroom.

5.3. The Post-game Debriefing Materials

As reasoned above, decision makers' mental activities on task knowledge acquisition and decision heuristics development will be captured from the post-game questionnaire. The whole debriefing session can be divided into two blocks. Subjects will be guided through each block of the questions without being allowed to refer back to the previous block in order to avoid mind-framing and self-justifying. The subjects will be given 20 minutes, half of the whole debriefing session, to fill in their answers for the open-ended question in the first block. There will be a reminder 5 minutes before the 20 minutes run out. Then they are guided to the next block containing closed-end questions for the rest of the 20 minutes. It will be stressed again that no communication between subjects and referring back to the previous block are allowed. Similarly, there will be a reminder 5 minutes before the end of the session.

The first block contains an open-ended question for heuristics development and task knowledge acquisition. The phrase will be: What information items were you looking at when you placed the weekly beer ordering? The subjects will be asked to list all of the information items they had employed in the game in terms of (1a) the variable's name, (1b) how the item was used to determine the weekly beer ordering, and (1c) how the item related to other variables that can be specified. The questions in (1b) and (1c) can be answered either qualitatively (such as "if the inventory was high, I would order less bear") or quantitatively (such as "I usually ordered beer as the difference between the current inventory and the weekly sales last time"). The question (1b) aims to extract decision heuristics developed; the question (1c), comparatively, is to extract task knowledge acquired.

There are a series of closed-end questions in Block 2 designed to measure whether there remain any task knowledge and decision heuristics developed during the gaming task which the subjects do not (or are not able to) explicate for the open-ended question in the previous block. There are two types of closed-ended questions for heuristics development: (2a) questions about the relations of the decision (the beer ordering) and performance measures (cost and inventory), (2b) questions about the relations between the decision and cues (e.g., weekly sales and arrivals). As the questions (2a) and (2b) relate the other task variables to the decision variable, the subjects may use these understanding to make decisions. Theoretically the answers for (2a) and (2b) may reflect the heuristics that may be potentially developed during the game playing.

Note that (2a) and (2b) may overlap with the subjects answers for the questions (1a) and (1b). That is why they are not allowed to go backward to revise the answers. Moreover, the answers for (2a) and (2b) also reflect the subjects' understanding about the task knowledge, which may be potentially utilized as decision rules. There is an additional type of questions to measure task knowledge acquisition: (2c) the questions about the relations between the cues and performance measures. The questions in (2c) may overlap with the open-ended questions in (1c). Specific questions for the debriefing session are still under development.

REFERENCES

Please download the full text of the article at
http://ALPHA1.albany.edu/~nh7365/prospectus.htm
for the detailed references.

APPENDIX. THE BEER DISTRIBUTION GAME

This appendix contains the following figures and tables that show the contents of the Beer Game. Refer to the web page of the System Dynamics Society at http://www.albany.edu/cpr/sds/ for more details.

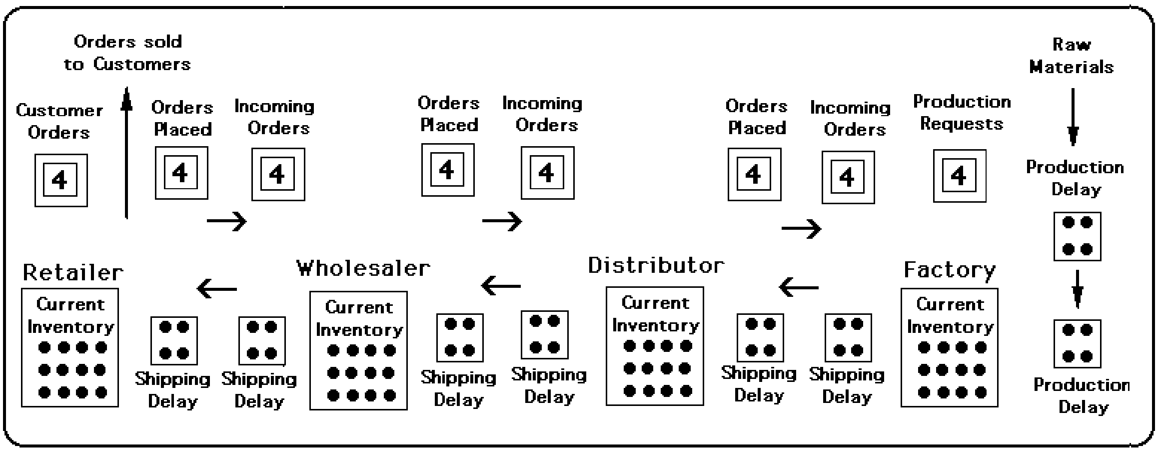

Figure B1: Game Board for the Original Beer Game