Carnegie Mellon University Research Showcase

Department of Social and Decision Sciences

Dietrich College of Humanities and Social Sciences

1-24-2005

Decision Support for Real-Time, Dynamic Decision-Making Tasks

Cleotilde Gonzalez
Carnegie Mellon University, conzalez@andrew.cmu.edu

Follow this and additional works at: http://repository.cmu.edu/sds

Recommended Citation

Gonzalez, Cleotilde, "Decision Support for Real-Time, Dynamic Decision-Making Tasks" (2005). Department of Social and Decision Sciences. Paper 11.

http://repository.cmu.edu/sds/11

This Article is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase. It has been accepted for inclusion in Department of Social and Decision Sciences by an authorized administrator of Research Showcase. For more information, please contact research-showcase@andrew.cmu.edu.





ORGANIZATIONAL
BEHAVIOR
AND HUMAN
DECISION PROCESSES

Organizational Behavior and Human Decision Processes 96 (2005) 142-154

www.elsevier.com/locate/obhdp

Decision support for real-time, dynamic decision-making tasks

Cleotilde Gonzalez*

Dynamic Decision Making Laboratory, Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Received 16 March 2004 Available online 24 January 2005

Abstract

By definition, dynamic decision making dictates that multiple and interrelated decisions be made in a continuously changing environment. Such decision making is difficult and often taxes individuals' cognitive resources. Here I investigated ways in which to support decision making in these environments. I evaluated three forms of decision support: outcome feedback, cognitive feedback, and feedforward that incorporated (to varying degrees) common features of learning theories associated with dynamic tasks. Participants in a laboratory experiment performed a real-time, dynamic decision-making task while receiving one of the different types of decision support. During the first 2 days, individuals received one type of decision support, but on the third day they performed the task without this support. Participants who received feedforward improved their performance considerably and continued to exhibit improved performance even after discontinuation of the decision support on the third day. Neither outcome feedback nor cognitive feedback resulted in improved performance. More research is necessary to conclusively identify the best forms of dynamic decision-making support and their durability when transferred to new tasks.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Dynamic decision making; Cognitive support; Feedback; Feedforward; Cognitive feedback

Introduction

Dynamic decision making (DDM) is characterized by the need to make multiple and interdependent decisions in an environment that changes as a function of the decision maker's actions, in response to environmental events, or in both ways (Edwards, 1962). Furthermore, because in dynamic environments decisions must be made in real-time, time constraints become an important determinant of performance (Brehmer, 1992). In addition, DDM is dynamically complex because it involves time delays and decisions that positively or negatively influence one another in complicated ways overtime (Diehl & Sterman, 1995).

Examples of DDM are: triage decisions in a medical emergency room; 911 operators determining relative

E-mail address: conzalez@andrew.cmu.edu.

urgency and deploying resources; drivers trying to find the best route in heavy traffic, military crew deciding how to respond to a threat or deploy forces, and air traffic controlling (ATC) attempting to allocate air space. ATC, for example, requires controllers to make multiple decisions regarding how to allocate space to best accommodate multiple airplanes. The fact that the assignment of a landing lane to an incoming airplane precludes the use of that lane by other airplanes arriving in the near future reflects the interdependency of decisions that characterizes DDM tasks. Furthermore, environmental parameters such as arrivals, departures, and weather are exogenous during ATC-i.e., they are beyond the influence of the controller. Finally, incoming airplanes need to be assigned to a landing lane at the correct moment in real-time. Thus ATC provides a realistic example of real-time DDM.

Because the successful performance of many important tasks requires skillful DDM, the identification of

^{*} Fax: +1 412 268 6938.

forms of decision support for DDM has become a research priority; however, this identification process has proven to be very challenging (Lerch & Harter, 2001). DDM researchers have investigated several types of decision support for DDM, including: (a) different types of outcome feedback that provide individuals with their performance results (Brehmer, 1995; Diehl & Sterman, 1995; Hsiao, 2000), (b) cognitive feedback that provides individuals with an explanation of how to perform the task (Sengupta & Abdel-Hamid, 1993), and (c) feedforward that helps decision makers examine the effects of possible future decisions (Lerch & Harter, 2001; Sengupta & Abdel-Hamid, 1993).

Prior research suggests that outcome feedback is an inefficient decision aid in situations involving complex and dynamic tasks (Hsiao, 2000; Lerch & Harter, 2001; Sterman, 1989) but it seems that the successful application of other forms of decision support (e.g., cognitive feedback or feedforward) may require frequent and timely outcome feedback (Balzer, Doherty, & O'Connor, 1989; Lerch & Harter, 2001).

Sengupta and Abdel-Hamid (1993) conducted a study in which they provided individuals performing a complex software management task with outcome feedback alone or in combination with either cognitive feedback or feedforward. Individuals assigned to the cognitive feedback group received graphs and tables containing detailed information about the key task variables and their interactions over time. Participants assigned to the feedforward group received a set of guidelines from individuals with experience in performing the task. In combination with outcome feedback, both cognitive feedback and feedforward resulted in better performance than did outcome feedback alone (Sengupta & Abdel-Hamid, 1993). Although this informative study yielded promising results, the task did not involve realtime decision making, which is a key characteristic of DDM tasks (Brehmer, 1992). Rather, individuals made decisions and processed feedback only at discrete time intervals. Furthermore, participants who received more detailed task information took longer and expended more cognitive effort while making their decisions than did participants who received only outcome feedback. The time constraints that are characteristic of real-time DDM tasks might make it impossible for individuals to efficiently process the type of feedback and feedforward provided in that study.

Lerch and Harter (2001) used a real-time DDM task to investigate the effects of outcome feedback and feedforward on performance. In their study, outcome feedback included explicit real-time (i.e., instantaneous) details about task performance. Feedforward involved a what-if computational analysis tool that allowed participants to 'look into the future' by observing the effects of possible actions. The results of that study indicate that the effectiveness of the support strategies depended on

the presence of outcome feedback. Feedforward alone impeded performance and inhibited learning, but feedforward provided in combination with outcome feedback led to slightly improved performance.

Because of the paucity of research focused on improving DDM, it is not entirely clear which forms of decision support might best aid individuals' performance of such tasks. Research indicates that feedback interventions may benefit individuals' performance of simple tasks but not of complex tasks, of tasks that demand low rather than high cognitive ability, and of well-practiced tasks but not novel ones (Kluger & DeNisi, 1996). These are discouraging results for researchers studying DDM support, because most DDM tasks are complex, demand high cognitive ability even after extensive task practice, and present decision makers with constantly changing and novel stimuli (Ackerman, 1992; Brehmer & Dörner, 1993; Gonzalez, Vanyukov, & Martin, 2005; Lerch & Harter, 2001)

The study reported here examined forms of decision support that appear to address the real-time constraints of DDM and the learning-process demands of DDM. It has been argued that achieving an improved understanding of individuals' learning processes will enable researchers to better understand performance in DDM tasks (and, subsequently, to design effective forms of decision support) (Gibson, Fichman, & Plaut, 1997; Gonzalez, Lerch, & Lebiere, 2003). The main DDM learning theories are discussed briefly below. Researchers also have established that individuals require extensive task practice to attain control over a dynamic system (Kerstholt & Raaijmakers, 1997) and have argued that the effectiveness of feedback interventions depends on task practice (Kluger & DeNisi, 1996). In the study presented here, participants working to complete a real-time DDM task were given frequent and detailed outcome feedback, cognitive feedback, or feedforward in an effort to better support learning in DDM. I examined the effects of these forms of decision support with extensive task practice and tested the permanence of performance improvements after practice by removing the decision support.

Learning in DDM

Although the psychology literature offers a large number of theories of learning, only a few of them are relevant to DDM. First, Dienes and Fahey (1995) proposed a model of learning based on two cognitive mechanisms that compete each time someone encounters a decision-making situation: an algorithm and a set of context-action exemplars (Dienes & Fahey, 1995; Logan, 1988). The algorithm is a general heuristic or rule that one uses in a novel situation; the context-action exemplars are discrete representations of knowledge that are called 'instances,' a name derived from Logan's

instance theory of automatization (1988). Implicit in this model is the assumption that a decision maker stores actions and their outcomes together in memory and retrieve them on the basis of their similarity to subsequently encountered situations. This model provides a feasible explanation of learning and performance in a dynamic task.

Second, Gibson et al. (1997) presented a model of learning in DDM based on the control theory proposed by Brehmer (1990) and implemented computationally via neural networks. In essence, this theory describes decision making in terms of two submodels: *The judgment submodel* learns by minimizing the differences between the outcomes it predicts and the actual outcomes, whereas the *choice submodel* learns by minimizing the differences between the choices predicted by the judgment model and the actual choices. This model provides a good account of individuals' learning in dynamic situations and their transfer of knowledge to novel situations (Gibson, 2000).

Third, Gonzalez et al. (2003) presented an instancebased model of learning in DDM (Instance-Based Learning Theory, IBLT). Borrowing from multiple theories of learning in psychology, IBLT proposes a learning process and a set of cognitive mechanisms. Not unlike Dienes and Fahey's theory, IBLT proposes that a decision maker acquires and stores instances in memory and retrieves them on the basis of their similarity to a subsequently encountered situation. In IBLT, however, these instances contain information about the utility of a decision in addition to information in the form of a situation-action pair. According to IBLT, learning occurs via the accumulation of instances initially stored with expected utility that are later revised on the basis of actual outcome feedback. As they accumulate these 'refined' instances, decision makers begin to recognize a particular situation as typical, and, in response to subsequent occurrences of similar situations, can retrieve from memory a decision with high utility. This model incorporates a foundational tenet of Logan's theory (1988): Judgment gradually moves from heuristic-based to recognition-based. The IBLT model was developed within ACT-R, a well-known theory of cognition (Anderson & Lebiere, 1998).

Although these three learning models differ in some respects, they all incorporate at least two common characteristics that may aid efforts to develop better decision support for DDM. First, all three models take into account the need for two forms of learning: explicit (i.e., decision making based on rules of action) and implicit (i.e., decision making based on context-based knowledge and recognition). There is some evidence that individuals who have completed a dynamic task are not always aware of the task structure (i.e., their knowledge is implicit), which suggests that the knowledge they acquired was not in the form of rules about how the sys-

tem works (Dienes & Fahey, 1995). Often, individuals performing DDM tasks are unable to describe the key elements of the task or verbalize the ways in which they make decisions (Berry & Broadbent, 1987, 1988). Such a lack of awareness both of the key variables involved in performing a task and of their relationships may denote an individual's dependence on implicit learning (Berry & Broadbent, 1987).

Second, these models, particularly the two instancebased models, rely on a similarity process that determines the applicability of accumulated experiences to familiar situations. Research in analogical reasoning has demonstrated the high relevance of analogy to learning and decision-making processes (Kurtz, Miao, & Gentner, 2001; Medin, Goldstone, & Markman, 1995). It appears that the *comparison* of tasks plays an important role in individuals' ability to identify common characteristics of tasks. For example, Loewenstein, Thompson, and Gentner (1999) have demonstrated that individuals who compared two analogous negotiation cases were likely to perform well in a subsequent negotiation task (Loewenstein et al., 1999). These results and others suggest that analogical encoding may facilitate abstraction and promote the transfer of knowledge (Thompson, Gentner, & Loewenstein, 2000).

The study presented here compared three forms of decision support that arise from these common features of the learning theories and from the characteristics of real-time DDM. To further investigate the effect of outcome feedback provided in conjunction with other forms of decision support, I provided some participants with frequent outcome feedback and global outcome feedback related to the decision-making task. As stated above, past research indicates that the provision of detailed and timely outcome feedback is a major factor in improving performance (Balzer et al., 1989). I hypothesized that frequent outcome feedback would help individuals to better determine the effects of their decisions, to improve the cause and effect association between their decisions and results, and, ultimately, to improve their decision-making performance.

Because the real-time nature of DDM imposes time constraints that limit decision makers' opportunities to search for and process information, decision makers must become aware of the task structure and the way they make decisions. Current learning theories suggest that people rely on a similarity process to make decisions while completing dynamic tasks. Analogical reasoning research suggests that comparing two tasks with different surface features but the same underlying structure helps individuals abstract and transfer knowledge. Therefore, I also tested the effect of an exemplar decision support in the form of cognitive feedback: I afforded individuals the opportunity to analyze their own past decisions without having to make decisions at the same time. This form of support

affords individuals the time they need to process information and takes advantage of some common learning features. I hypothesized that individuals permitted to review their own past decisions in the absence of time constraints would analyze their actions, improve their decisions, and exhibit improvements in overall task performance.

I also tested another exemplar decision support in the form of feedforward: I afforded individuals the opportunity to compare their decisions with those made by an expert performer. Research suggests that individuals may be able to improve their performance by comparing their decisions and effects of their decisions with the decisions made by an expert and the effects of those decisions (Sengupta & Abdel-Hamid, 1993). Because it allows individuals to analyze expert's decisions without having to execute decisions at the same time, this feedforward support removes time constraints. I hypothesized that individuals permitted to review an expert's decisions without time constraints would exhibit improvements in overall task performance.

Methods

Water purification plant: A DDM task

The research reported here involved a computer simulation called 'Water Purification Plant' (WPP). WPP is a computer-based, interactive, simulation isomorph of a real-world scheduling task: mail sorting in the United States Postal Service (USPS). An in-depth description of the original USPS task is available elsewhere (Lerch, Ballou, & Harter, 1997). As reported in previous research (Gonzalez, 2004; Gonzalez et al., 2003), WPP incorporates the parameters involved in the original USPS task, but WPP features a simplified interface that helps in both explaining the task and learning it more quickly. The main elements of WPP are summarized below.

WPP simulates a water distribution system with 23 tanks (tanks 0–22) arranged in a treelike structure and connected with pipes. The goal in this task is to distribute all the water in the system before the various deadlines expire. The tanks, pipe connections, and deadlines are shown in the screenshot of the simulation in Fig. 1. Each tank has two pumps, and a total of only five pumps can be active within the entire system at any given time. (Participants are told that electricity constraints preclude the concurrent use of more than five pumps.) By activating different pumps, users attempt to distribute water through the pipes before deadlines expire. Each

trial covers a time span of 8 simulation hours (2:00–10:00). Each tank's fixed position within the tree structure dictates the number of tanks the water must travel through before expiration of the deadline. Each pump delivers water at a rate of 1 gal every 2 min; thus when two pumps in one tank are active, the delivery rate is 1 gal/min. The pattern according to which water enters the system is defined by a preset scenario that dictates the arrival time, the amount of water, and the destination tank. For example, an entry in the WPP scenario may indicate that at 2:02 10 gal water will enter tank 2. The water entry scenario is unknown to users.

The simulation provides an indicator that tracks the number of pumps in use (shown at the top left corner in Fig. 1). Each pump may be in one of four states, which are identifiable by the color of the indicator bar above each pump: off (red), on (green), cleaning (yellow), or in queue (purple). After being deactivated (either by the operator or by the system when there is no more water remaining in the tank), a pump switches into cleaning mode and remains in that state for 10 min of simulation time. While pumps are cleaning, the user can select other pumps (within the 5 pump limit) that will become active as soon as the cleaning pump turns red. Pumps so identified by the user are in queue (purple).

WPP embodies the integral characteristics of all DDM tasks (Brehmer & Dörner, 1993; Gonzalez et al., 2005). First, WPP is a dynamic resource allocation task. In WPP, exogenous events (i.e., the arrival of water at different times throughout the simulation) harm the status of the system (e.g., remain in the chains after deadline expiration) if no action is taken. The users' actions, which are restricted by the limited resources, also affect the status of the system. Second, WPP is a complex task because it involves multiple variables (e.g., pumps, water, and deadlines) and because some of the relationships among these variables are non-linear. For example, pumping water out of a tank before the deadline reduces the number of gallons missed, but pumping water out of a tank after the deadline has expired has no effect on performance. Third, WPP is opaque in the sense that many characteristics of the system are not visually discernable and are identifiable only by user inference. For example, later deadlines have longer chains, and more water may accumulate in the tanks. This situation may redefine the task prioritization suggested by the deadlines. In addition, WPP is dynamically complex. Decisions' effects on other variables over time result in feedback loops (Sterman, 1989). For example, the more often pumps are activated the more time they spend in cleaning mode and the less time is available for pumping water through the chain.

Performance variable

Performance in WPP is measured by the total number of gallons of water that remain in chains after the

¹ To make the figure more clear, I added explanations regarding the different parts of the figure and I numbered the tanks. This information does not appear during the simulation.

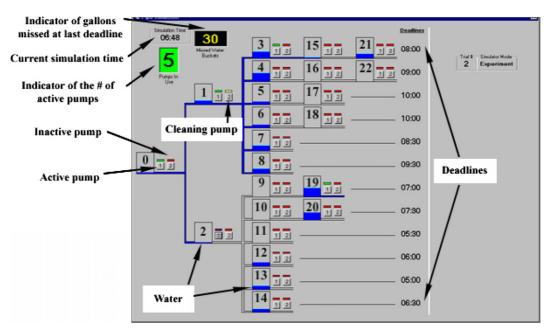


Fig. 1. Screenshot of the WPP simulation. *Note*. The simulation time is 6:48, the operator has missed 30 gal water, and 5 pumps (the maximum number possible) are active. Water enters from outside the system and moves continuously through the open pumps from left to right toward the deadline column. The operator decides when to open and close pumps while simulation time is running. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this paper.)

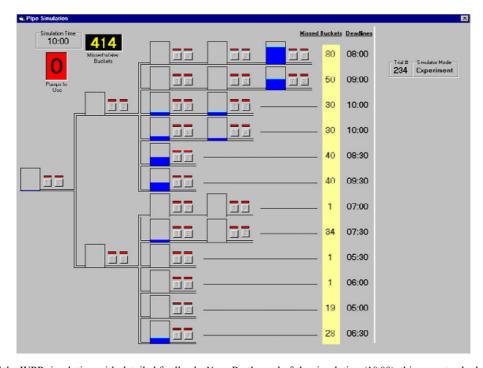


Fig. 2. Screenshot of the WPP simulation with detailed feedback. *Note*. By the end of the simulation (10.00), this operator had missed 414 gal water. The number of gallons missed per deadline was presented next to each of the deadlines as they expired during the simulation. (The numbers were highlighted for presentation purposes only.)

expiration of all deadlines. The simulation also presents participants with a cumulative number of gallons of water remaining in the system (the score appears in the top left corner of the screen, as shown in Fig. 1). The system's total capacity per trial is 1080 gal water. The score from the WPP simulation is used to calculate a perfor-

mance percentage measure. Larger percentages (fewer remaining gallons) indicate better performance, and the highest possible performance is 100% (1080 gal pumped/0 gal remaining).

In this task, as in many other complex tasks, optimal performance is achievable via many possible decision

Table 1 Experimental design

Condition	Practice		Test		
	Days 1 and 2	Day 3			
	Odd trials (1, 3,, 11)	Even trials (2,4,,12)	Trials 13–18		
Control	Execution of 6 trials at 8 min per trial	Execution of 6 trials at 8 min per trial	Execution of 6 trials at 8 min per trial		
Feedback	Execution of 6 trials with frequent feedback at 8 min/trial	Execution of 6 trials with frequent feedback at 8 min/trial	Execution of 6 trials at 8 min/trial		
Self-exemplar	Execution of 6 trials at 8 min/trial	8 min <i>replay of the previous</i> odd-numbered trial	Execution of 6 trials at 8 min/trial		
Feedback- exemplar	Execution of 6 trials with frequent feedback at 8 min/trial	8 min <i>replay of the previous</i> odd-numbered trial	Execution of 6 trials at 8 min/trial		
Expert- exemplar	Execution of 6 trials at 8 min per trial	8 min replay of an expert's trial	Execution of 6 trials at 8 min per trial		

sequences. To establish a reasonable lower limit for task performance, my colleagues and I created a program to run the simulation by activating pumps randomly while allowing no idle time (i.e., always having five active pumps). Thirty replications produced a mean performance percentage of 81.9% with a standard deviation of 2.4%. Therefore, for the purposes of this study, a performance percentage around 80% was considered a reasonable first-time performance, and, with practice and learning, participants were expected to attain performance percentages above 80%. Over the past few years, approximately 500 participants have run WPP with practice times ranging from 3 to 20h. The best past performance percentages have ranged from 90 to 95%; no one has achieved the best possible performance percentage (i.e., 100%).

Experimental design

Table 1 outlines the key elements of the experimental design. Participants were assigned to one of five conditions. Under the *control* condition, participants ran WPP 18 times over the course of 3 days (6 trials/day, 8 min/trial). The only feedback provided to individuals in this group was the total number of gallons remaining in the system after expiration of all deadlines (i.e., the users in this group received no feedback until the end of each trial).

Individuals in the *feedback* group received detailed performance-outcome feedback in addition to the outcome feedback described above. The more detailed feedback included the number of gallons of water remaining in the system upon expiration of each of the deadlines on the screen, rather than just the accumulated number of gallons at the end of the trial (as provided under the control condition; see Fig. 2). These participants ran WPP 12 times during the first 2 days (6 trials/day, 8 min/trial), and then ran 6 trials under the control condition (see above) on the third day.

Participants in the *self-exemplar* group ran one trial under the control condition, and then replayed that trial.

The WPP replay function replays each action taken in the original trial and exactly reproduces every aspect of the original screen as the task replays. Individuals were instructed to observe and analyze their decisions as they replayed the original trial. The participants in this group followed this procedure for the first 2 days by running trials 1, 3, 5, 7, 9, and 11. Trials 2, 4, 6, 8, 10, and 12 were replays of the original set of trials (i.e., the odd-numbered trials). Each replay lasted as long as the original task simulation (i.e., 8 min/trial). Viewing the replays was expected to afford participants time to better encode and re-evaluate previously made decisions because the participants did not need to interact with the simulation while they analyzed their past actions. On the third day, participants in this group ran six trials under the control condition.

Participants in *the feedback-exemplar* group received both detailed feedback and the opportunity to replay their own trials as described above. During the first 2 days, these participants were given detailed feedback while interacting with WPP and were permitted to replay their trials as described above. On the third day, these individuals ran six trials of WPP under the control condition.

The expert-exemplar condition was similar to the self-exemplar condition, except that individuals replayed the trial of a highly skilled participant (achieving the best possible performance) rather than their own previous trials. Viewing replays of an expert's trial was expected to give participants the time needed to encode and evaluate the decisions of an expert player, and presumably to help them improve their own performance. On the third day, the expert-exemplar participants ran six trials under the control condition.

Participants

Eighty-eight graduate and undergraduate college students recruited from local universities were randomly assigned to one of the five conditions (14 to control, 19 to feedback, 16 to self-exemplar, 20 to feedback-exemplar, and 19 to expert-exemplar). The average age of the participants was 21 years; 48% of the participants were female and 52% were male. Each individual received \$50 at the end of the three-consecutive-day study.

Procedure

On the first day, participants received instructions outlining the goals of WPP and how to use the simulation. Instructions were given according to a standard script while the participants ran the simulation in a training mode (at the very slow pace and only until expiration of the first deadline). During the instructional period, participants were informed of deadlines, simulation time, and the routes of water travel but were not instructed or given any information on how to improve their performance. They were told that water (from outside the system) could arrive in any of the tanks at any time, but they were not given information about the amount of water to process or the time of water arrival.

Statistics

I used repeated measures ANOVAs and one-way ANOVAs to analyze the data. Because the even-numbered trials (2, 4, 6, 8, 10, and 12) were replays in the self-exemplar, feedback-exemplar, and expert-exemplar conditions, no performance data from those groups were available for those trials. Therefore, those trials were dropped from the analyses of all groups. For this reason, trials referred to as trials 1–6 in the Results and Discussion below correspond to actual performance in the odd-numbered trials (1, 3, 5, 7, 9, and 11), whereas trials referred to as trials 7–12 in the remainder of the paper

correspond to the test phase (day 3) trials 13–18. For the repeated measures ANOVAs, the condition (control, feedback, self-exemplar, feedback-exemplar, or expert-exemplar) was the between-subjects factor and the trials (a total of 12 trials, six during the first 2 days, and six during the third day) were the repeated measures of performance. Significance was assigned at the level of p < .05.

Results

Table 2 presents the mean performance percentages and standard deviations for each of the conditions during the 12 trials; Fig. 3 depicts the mean performance percentages graphically. The overall analysis revealed a significant main effect of trials (F(11, 913) = 78.91, p < .001), a trials-by-condition significant interaction (F(44, 913) = 2.43, p < .001), and a significant main effect of condition (F(4, 83) = 2.81, p < .05). Thus, individuals tended to improve their performance over time and to perform differently according to the availability and type of cognitive support. Performance improvements differed on the basis of condition.

Comparisons between each of the experimental groups and the control group (across trials) produced the results summarized in Table 3. Individuals in the expert-exemplar group were the only ones to perform at significantly higher levels than individuals in the control group (F(1, 31) = 5.96, p = .021). Results in the feedback, self-exemplar, and feedback-exemplar groups did not differ significantly from those in the control group. The analyses also revealed a significant interaction of trials-by-condition for the feedback group (F(11, 341) = 3.07, p < .005) and the expert-exemplar group (F(11, 341) = 3.72, p < .001) as compared (respectively) with the control group.

Table 2
Mean performance percentages, standard errors of means, and standard deviations for each condition during the 12 trials

Condition		Days 1 and 2						Day 3					
		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	Trial 11	Trial 12
Control $(N = 14)$	M	80.67	84.89	86.00	86.04	87.59	88.60	88.44	89.96	88.88	88.42	89.70	89.62
	σ_{M}	2.09	1.60	1.10	1.32	1.50	1.01	1.06	1.14	1.30	1.60	1.53	1.30
	SD	7.82	5.98	4.11	4.96	5.63	3.77	3.95	4.28	4.86	5.97	5.73	4.88
Feedback $(N = 19)$	M	74.73	76.54	82.71	84.57	87.41	87.89	89.03	89.27	89.29	88.84	89.85	90.07
	σ_M	1.72	3.78	3.05	2.37	1.27	1.35	1.42	1.30	1.57	1.60	1.61	1.52
	SD	7.49	16.47	13.31	10.32	5.54	5.88	6.18	5.65	6.84	6.98	7.03	6.61
Self-exemplar $(N = 16)$	M	81.02	84.53	87.15	89.41	89.09	90.07	89.22	90.72	90.20	90.69	89.84	90.54
	σ_M	1.68	1.06	1.06	1.13	1.14	1.08	1.30	1.12	1.09	1.34	1.40	1.24
	SD	6.73	4.23	4.22	4.51	4.58	4.31	5.19	4.47	4.36	5.34	5.60	4.95
Feedback-exemplar $(N = 20)$	M	77.25	83.50	85.62	86.43	87.05	88.16	88.94	88.81	90.23	90.60	91.11	91.00
• • • • • • • • • • • • • • • • • • • •	σ_M	1.55	1.25	1.31	1.12	1.20	1.19	1.30	1.83	1.29	1.04	1.01	0.95
	SD	6.95	5.61	5.88	5.02	5.38	5.32	5.82	8.18	5.78	4.67	4.52	4.26
Expert-exemplar $(N = 19)$	M	77.13	88.71	88.76	89.98	92.79	93.15	94.10	94.15	93.05	93.19	93.50	93.78
	σ_M	2.39	1.51	1.54	1.09	0.94	1.01	0.98	0.70	0.98	0.10	1.00	1.00
	SD	10.41	6.58	6.70	4.73	4.11	4.39	4.26	3.06	4.25	4.31	4.35	4.30

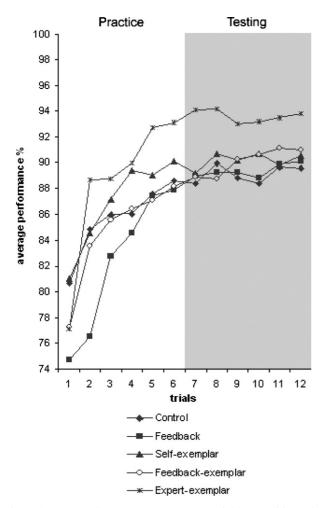


Fig. 3. Average performance percentage per trial by condition. The white area of the graph represents the practice period of each condition, and the gray area represents the testing period of each condition.

Table 3
Planned comparisons of each condition versus the control condition

Condition	F	p
Feedback vs. Control		
Trial effect	(11, 341) = 22.8	0.00
Trial*Condition	(11, 341) = 3.07	0.005
Condition	(1,31) = 0.52	n.s.
Self-exemplar vs. Control		
Trial effect	(11, 308) = 18.7	0.00
Trial*Condition	(11, 308) = 0.59	n.s.
Condition	(1, 28) = 0.66	n.s.
Feedback-exemplar vs. Co	ntrol	
Trial Effect	(11, 352) = 27.7	0.00
Trial*Condition	(11, 352) = 1.55	n.s.
Condition	(1, 32) = 0.00	n.s.
Expert-exemplar vs. Contro	ol	
Trial effect	(11, 341) = 36.2	0.00
Trial*Condition	(11, 341) = 3.72	0.001
Condition	(1,31) = 5.96	0.021

To further analyze these two trials-by-condition interactions, I ran two separate analyses on the data: one analysis (using the six trials as a repeated measures

variable) for the first 2 days, during which individuals received different degrees or types of feedback or replays, and the other for the third day when all participants were under the control condition. These analyses indicated that the performance of individuals in the feedback group differed from the performance of those in the control group during the first 2 days (F(5,155) = 2.61, p < .05) but not during the third day (F(5, -1)) 155) = .25, p > .05). Per-trial analyses showed that the difference during the first 2 days occurred only during the first trial (F(1, 32) = 4.89, p < .05). Initially, individuals given detailed feedback performed significantly more poorly than individuals in the control group, but with practice the individuals who received detailed feedback performed at levels comparable to those exhibited by individuals in the control group. When individuals in the detailed-feedback group completed the simulation under the control condition on the third day of the study, they performed exactly like the individuals in the control group, who had never received any detailed feedback.

Per-trial analyses performed to investigate the learning differences between individuals in the expert-exemplar group and those in the control group revealed that the former group of participants required only three task executions (and two replays of an expert's trial) to outperform participants in the control group; from the fourth trial forward, individuals in the expert-exemplar group performed better than individuals in the control group. Table 4 summarizes the analyses per trial. Performance in the expert-exemplar group differed significantly from that in the control group on the third day (F(1, 31) = 9.64, p < .005), and the performance of the expert-exemplar participants continued to improve even in the absence of the replay of an expert's trial (i.e., on the third day).

In summary, only participants under the expert-exemplar condition performed better than participants under the control condition. The feedback condition actually led to poorer performance at the beginning of the learning process, and the exemplar and feedback-exemplar conditions had no significant effect on performance.

Process variables

To determine why performance under the experimental conditions did or did not differ from that under the control condition, I analyzed two process variables: the percentage of pump utilization and the average number of decisions per trial. The percentage of pump utilization was calculated by dividing the total number of minutes that pumps were open (480) by the total number of minutes of pump availability (480 min [the total duration of the simulation] multiplied by 5 [the maximum number of pumps open at any one time]). The second process variable was the number of decisions made by each

Table 4 Planned comparisons: one-way ANOVAs per trial

Condition	Days 1 a	and 2	•	•	•		Day 3					
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	Trial 11	Trial 12
Feedback vs. F(1, 31)	Control 4.89**	3.25*	0.79	0.24	0.01	0.16	0.10	0.15	0.04	0.03	0.01	0.05
Self-exemple $F(1, 28)$	ar vs. Contr 0.02	ol 0.04	0.57	3.81*	0.64	0.97	0.21	0.23	0.61	1.21	0.01	0.26
Feedback-ex F(1, 32)	emplar vs. 1.79	Control 0.48	0.04	0.05	0.08	0.07	0.07	0.23	0.51	1.42	0.64	0.78
Expert-exem F(1, 31)	ıplar vs. Ca 1.14	ontrol 2.94*	1.85	5.36**	9.42***	9.72***	15.08***	10.78***	6. 87**	7.12**	4.73**	6.75**

^{*} $p \le .10$.

individual. In WPP, the time-related cost of switching pumps (i.e., 10 simulation minutes of inactivity due to pump cleaning) should discourage the activation of pumps without careful consideration. Thus the number of decisions per trial also serves as a measure of efficient use of resources. The higher the number of decisions per trial, the more the decision maker changed her decisions regarding pump activation, potentially resulting in negative effects on performance. As a raw measure of efficiency, the expert's trial used for this study involved 29 decisions.

Using these two variables, I ran regression analyses on the average performance during the first 2 days and during the third day. Both regressions yielded significant results (days 1 and 2, Adjusted $R^2 = .85$, F(2, 87) = 174.05, p < .001; day 3, Adjusted $R^2 = .78$, F(2, 87) = 154.37, p < .001). Regression coefficients generated by both analyses also were significant for the percentage of pump utilization and the number of decisions per trial during the first 2 days and during the third day. These results suggest that the differences in performance between individuals in the control group and those in the other groups are explainable in terms of the percentage of pump utilization and the average number of decisions per trial.

I also used repeated measures analyses to compare the percentage of pump utilization by individuals in each of the experimental groups to the percentage of pump utilization by participants in the control group. Like the performance analyses, these analyses revealed that only the individuals in the expert-exemplar group utilized the pumps better than individuals in the control group. Although the difference between the control and expert-exemplar groups was not significant during the first 2 days (F(1, 31) = .079, p = .78), it was significant during the third day (F(1, 31) = 5.57, p < .05). Pump utilization by individuals in the feedback, exemplar, or feedback-exemplar groups. Fig. 4 shows the

average percentage of pump utilization over the course of the study by individuals in the expert-exemplar and control groups. The expert's trial replayed for this study involved 98% pump utilization.

As with the percentage of pump utilization analyses, repeated measures analyses using number of decisions

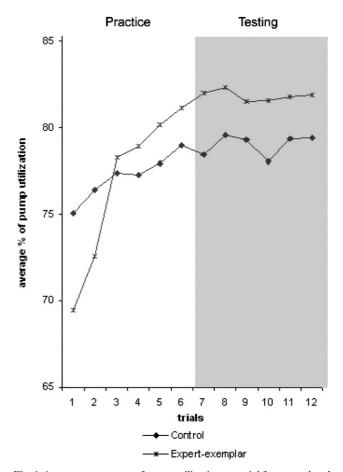


Fig. 4. Average percentage of pump utilization per trial for control and expert-exemplar conditions. The white area of the graph represents the practice period of each condition, and the gray area represents the testing period of each condition.

^{**} $p \le .05$.

 $p \leq .01$.

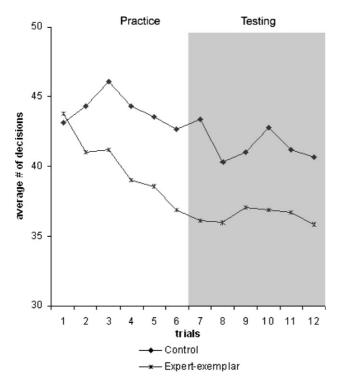


Fig. 5. Average number of decisions, over the course of the study, for the expert-exemplar and control conditions. The white area of the graph represents the practice period of each condition, and the gray area represents the testing period of each condition.

per trial were run to compare each experimental group to the control group. Once again, the only condition significantly different from the control condition was the expert-exemplar condition (F(1, 31) = 15.66, p < .001). The number of decisions made by individuals in the control group did not differ from the number made by individuals in the feedback, exemplar, or feedback-exemplar groups. Fig. 5 shows the average number of decisions made over the course of the study by individuals in the expert-exemplar and control groups.

Results from the analyses described above indicate that individuals who replayed the expert's trial improved their resource utilization and made a reduced number of decisions after practicing the task. With practice, these individuals performed at levels that increasingly approached the expert's performance. Presumably, these participants compared their decisions with those made by the expert and, in subsequent trials, attempted to duplicate the expert's decision-making approaches.

Similarity

To test whether the decisions made by individuals in the expert-exemplar group were, in fact, increasingly similar to those made by the expert performer, I ran an analysis of decision similarity by comparing each of the individuals' trials to the expert's trial. The measure of similarity used here, which was proposed and tested in previous research (Gonzalez et al., 2003; Gonzalez & Quesada, 2003), is the linear distance between two decisions, and it ranges from 0 (diametrically dissimilar) to 1 (equivalent). In contrast with the previous analysis, however, this analysis compared the decisions of each individual to those of the expert rather than to each individual's past decisions.

The analysis of similarity showed a very significant, main effect of the condition (F(4, 83) = 7.96, p < .001). The systematic comparisons of each experimental group to the control group showed a significant difference between the control and the expert-exemplar groups (F(1, 31) = 16.87, p < .001). Again, the similarity between the expert's decisions and the decisions made by individuals in the feedback, self-exemplar, and feedback-exemplar groups did not differ from the similarity between the expert's decisions and the decisions made by individuals in the control group. Fig. 6 shows the average similarity between the decisions made by individuals in the control and expert-exemplar groups and those made by the expert.

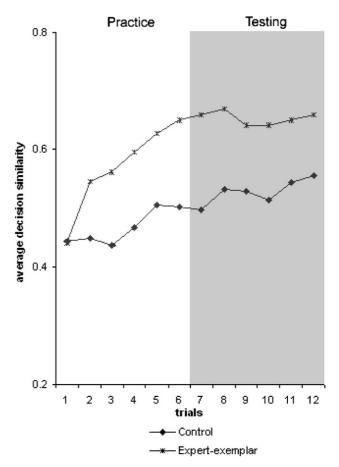


Fig. 6. Average similarity of the decisions made in the control and the expert-exemplar conditions to the decisions made in the expert's (replayed) trial. The white area of the graph represents the practice period of each condition, and the gray area represents the testing period of each condition.

Discussion

The results of this study are consistent with prior research results that have shown that outcome feedback alone is an ineffective form of decision support. This study also demonstrates that more detailed and frequent outcome feedback did not help individuals improve their performance. Furthermore, process feedback in the form of a replay of individuals' own past performance was also ineffective. Only one feedforward decision aid was successful: the expert-exemplar condition that involved replaying the decisions made by an expert. It is interesting that individuals continued to improve their performance even after cessation of this feedforward decision support.

The results of this study resonate with prior research results that have shown that feedforward in the form of information from experienced individuals can help improve performance (Sengupta & Abdel-Hamid, 1993). Furthermore, this study demonstrates a decision aid that is useful as support for *real-time* DDM. Rather than asking individuals to processing feedforward in the form of guidelines while they perform the task, I asked participants to replay an expert's trial after they had performed the task. This enabled participants to compare their own decisions with those of an expert, and thereby to improve their own performance. This type of decision support is unobtrusive because it corresponds to the natural form of learning in DDM tasks.

Individuals did not need to be consciously aware of the heuristics or processes that the expert followed; rather, individuals may simply have used examples as a natural form of learning. As suggested by the instance-based learning theories, individuals may have learned by accumulating instances from both their own trial and the expert's trial (Dienes & Fahey, 1995; Gonzalez et al., 2003). Participants may have stored high utility instances while playing the expert's trial, and then retrieved those instances when playing their own trials. The similarity analysis indicates that, unlike the decisions made by individuals in any of the other groups, the decisions made by individuals in the expert-exemplar group became increasingly similar to those made by the expert.

Although it is possible that decision makers simply copied the expert's behavior analyses of the process variables suggest that reviewing the expert's decisions helped individuals manage their own resources more wisely. Individuals learned to utilize the pumps more effectively while making fewer decisions. More importantly, individuals' utilization of resources continued to improve even after discontinuation of the expert-exemplar decision support. There are several possible explanations for the performance improvements in the absence of decision support. First, it is feasible that participants, rather than simply copying the expert's behavior, were able to

understand the task beyond a particular example and to construct a more general strategy that aided them in solving the task. Second, viewing the expert's decisions may have helped participants retrieve actions in the neighborhood of the expert's decisions, giving participants a general sense of strategy that they could then fine-tune with practice. That is, if instances of the expert's performance become part of the participants' memory repertoire, participants may have used that knowledge to make decisions that led to higher performance even in the absence of the decision support. Additional empirical research should be performed to assess the validity of these two possible explanations.

DDM researchers understand very little about what people learn while completing DDM tasks and what components of learning are transferable. Past studies have shown that, although individuals may improve their performance in a specific task, they are unable to transfer this knowledge even to very similar tasks (Sterman, 1989). To improve our understanding regarding the effectiveness and durability of decision support, researchers must conduct additional transfer-of-learning studies that test individuals' ability to transfer learning when confronted with different but analogous conditions.

It is also important to consider why the other forms of decision support investigated in this study failed. As predicted by Kluger and DeNisi's theory, this study showed that outcome feedback impeded performance at the beginning of the practice period (when the task was novel) but had fewer detrimental effects on the performance of individuals after additional task practice. This finding provides strong evidence of the ineffectiveness of outcome feedback, regardless of specificity. Contrary to expectations, the timing and frequency of outcome feedback did not help improve performance.

There are several possible explanations for this finding, and some of them have already been addressed by other researchers (Sengupta & Abdel-Hamid, 1993). One possible explanation has to do with individuals' limited cognitive resources. Because real-time dynamic environments do not pause to allow for feedback presentation, they force decision makers to continue making judgments while a situation progresses. Processing feedback while determining how to alter one's judgments according to that feedback is a difficult and time-consuming task (Einhorn & Hogarth, 1978). The inefficient utilization of resources during the early trials under this condition—which at least partially explains the lower initial performance of individuals given feedback—may be attributable to the participants' inability to perform both tasks concurrently. Perhaps individuals attempted to use the feedback during the first few trials but later simply ignored it because they could not effectively process the feedback while continuing to perform well on the WPP task.

The failure of the self-exemplar decision support could be attributed to the quality of information upon which individuals relied during the trials. IBLT suggests that participants in this group may have used their own prior decisions to formulate instances; however, the utility of these instances was likely low. Participants did not seem to have learned anything new from their past decisions. Research suggests that most people do not account very well for the impact of their own decisions (Sterman, 1989). According to Einhorn and Hogarth (1978), the difficulty of learning from experience is attributable to three factors: (1) the failure to search for and use disconfirming evidence, (2) a lack of awareness of environmental effects on outcomes, and (3) the use of unaided memory for encoding, storing, and retrieving outcome information. Although it is possible that individuals in the self-exemplar group had trouble encoding, storing, and retrieving past decisions, this seems rather unlikely, particularly upon comparison of the data from this group with those from the expert-exemplar group. Individuals in the expert-exemplar group had no trouble encoding, storing, and retrieving information similar to that of the expert. It seems more likely that, in this case, the first two factors may have played a role in the ineffectiveness of the self-exemplar decision aid. Replaying their own decisions may not have helped individuals to understand the effects of their decisions, and individuals may have had little motivation to seek alternate courses of action.

Results from this study confirm the ineffectiveness of using outcome feedback alone to support decision making. Findings also indicate that the provision of more frequent and specific outcome feedback (rather than only global outcome feedback) hinders rather than supports performance. Furthermore, access to cognitive feedback based on the individuals' replays of their own trials did not improve performance. This finding indicates that affording individuals the time and opportunity to analyze their actions is not enough to help them improve performance.

However, this study also demonstrates a form of decision support that helped individuals improve their performance considerably: the replay of an expert's trial. This result is applicable to many real-world dynamic tasks because, in the real world, there are well-recognized experts in many different fields. Often, knowledge is sought from experts to examine decision making. However, the results of this study suggest that it might be easier and more effective to solicit examples of experts' decisions and use them as decision support for trainees. The use of realistic simulations that reproduce real-world tasks—called "management flight simulators" by Sterman (Sterman, n.d.)—may encourage analogical reasoning and provide a means of extensive practice by individuals who must transfer the knowledge gained through such simulations to real-world tasks.

Sterman argues that such simulations have helped to improve productivity and quality in many organizations.

Acknowledgments

This research was supported by the Multidisciplinary University Research Initiative Program (MURI; N00014-01-1-0677) and by the Advanced Decision Architectures Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory (DAAD19-01-2-0009). I am also grateful for editorial assistance provided by Ryan Sauder during preparation of the manuscript.

References

- Ackerman, P. L. (1992). Predicting individual differences in complex skill acquisition: Dynamics of ability determinant. *Journal of Applied Psychology*, 77(5), 598–614.
- Anderson, J. R., & Lebiere, C. (1998). The atomic components of thought. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Balzer, W. K., Doherty, M. E., & O'Connor, R. (1989). Effects of cognitive feedback on performance. *Psychological Bulletin*, 106(3), 410–433.
- Berry, D. C., & Broadbent, D. E. (1987). The combination of explicit and implicit learning processes in task control. *Psychological Research*, 49(1), 7–15.
- Berry, D. C., & Broadbent, D. E. (1988). On the relationship between task performance and associated task knowledge. *Quarterly Journal of Experimental Psychology A*, 36, 209–231.
- Brehmer, B. (1990). Strategies in real-time, dynamic decision making. In R. M. Hogarth (Ed.), *Insights in decision making* (pp. 262–279). Chicago, IL: University of Chicago Press.
- Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologies*, 81(3), 211–241.
- Brehmer, B. (1995). Feedback delays in complex dynamic decision tasks.
 In P. A. Frensch & J. Funke (Eds.), Complex problem solving: The European perspective. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Brehmer, B., & Dörner, D. (1993). Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior*, 9(2–3), 171–184.
- Diehl, E., & Sterman, J. D. (1995). Effects of feedback complexity on dynamic decision making. Organizational Behavior and Human Decision Processes, 62(2), 198–215.
- Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 21(4), 848–862.
- Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human factors*, 4, 59–73.
- Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgement: Persistence of the illusion of validity. *Psychological Review*, 85(5), 395–416.
- Gibson, F. P. (2000). Feedback delays: How can decision makers learn not to buy a new car every time the garage is empty. *Organizational Behavior & Human Decision Processes*, 83(1), 141–166.
- Gibson, F. P., Fichman, M., & Plaut, D. C. (1997). Learning in dynamic decision tasks: Computational model and empirical evidence. *Organizational Behavior and Human Decision Processes*, 71, 1–35.
- Gonzalez, C. (2004). Learning to make decisions in dynamic environments: Effects of time constraints and cognitive abilities. *Human Factors*, 46(3), 449–460.

- Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. Cognitive Science, 27(4), 591–635.
- Gonzalez, C., & Quesada, J. (2003). Learning in dynamic decision making: The recognition process. *Computational & Mathematical Organization Theory*, 9(4), 287–304.
- Gonzalez, C., Vanyukov, P., & Martin, M. K. (2005). The use of microworlds to study dynamic decision making. *Computers in Human Behavior*, 21, 273–286.
- Hsiao, N. (2000). Exploration of outcome feedback for dynamic decision making. University at Albany, State University of New York, Albany, NY..
- Kerstholt, J. H., & Raaijmakers, J. G. W. (1997). Decision making in dynamic task environments. In R. Ranyard, R. W. Crozier, & O. Svenson (Eds.), *Decision making: Cognitive models and explana*tions. Norwood, NJ: Ablex.
- Kluger, A. N., & DeNisi, A. (1996). Effects of feedback intervention on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin*, 119(2), 254–284.
- Kurtz, K. J., Miao, C., & Gentner, D. (2001). Learning by analogical bootstrapping. The Journal of Learning Sciences, 10(4), 417–446.
- Lerch, F. J., Ballou, D. B., & Harter, D. E. (1997). Using simulation-based experiments for software requirements engineering. *Annals of Software Engineering*, 3, 345–366.

- Lerch, F. J., & Harter, D. E. (2001). Cognitive support for real-time dynamic decision making. *Information Systems Research*, 12(1), 63– 82
- Loewenstein, J., Thompson, W. L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psycho-nomic Bulletin and Review*, 6(4), 586–597.
- Logan, G. (1988). Toward an instance theory of automatization. *Psychological Review*, 95(4), 492–527.
- Medin, D. L., Goldstone, R. L., & Markman, A. B. (1995). Comparison and choice: Relations between similarity processing and decision processing. *Psychonomic Bulletin and Review*, 2(1), 1–19.
- Sengupta, K., & Abdel-Hamid (1993). Alternative conceptions of feed-back in dynamic decision environments: An experimental investigation. *Management Science*, 39(4), 411–428.
- Sterman, J. (1989). Misperceptions of feedback in dynamic decision making. Organizational Behavior and Human Decision Processes, 43(3), 301–335.
- Sterman, J. (n.d.). Teaching takes off: Flight simulators for management education: "The Beer Game." Available from: http://web.mit.edu/jsterman/www/SDG/beergame.html.
- Thompson, W. L., Gentner, D., & Loewenstein, J. (2000). Avoiding missing opportunities in managerial life: Analogical training more powerful than individual case training. Organizational Behavior & Human Decision Processes, 82(1), 60–75.