

Dynamic Decision Making

Jared M Hotelling, University of Basel, Basel, Switzerland

Pegah Fakhari and Jerome R Busemeyer, Indiana University, Bloomington, IN, USA

© 2015 Elsevier Ltd. All rights reserved.

This article is a revision of the previous edition article by J.R. Busemeyer, vol 6, pp. 3903–3908, © 2001, Elsevier Ltd.

Abstract

This section reviews a specialty within the field of decision making known as dynamic decision making. Dynamic decisions are characterized by a decision maker choosing among various actions at different points in time in order to control and optimize performance of a dynamic stochastic system. Realistic examples include fighting fires, navigational control, battlefield decisions, medical emergencies, and so on. The section has four parts. The first reviews basic theory concerning optimal decision principles in a dynamic context, the second summarizes empirical approaches to the study of human performance on dynamic decision tasks, the third presents theoretical models that describe how humans learn to control dynamic systems, and the last discusses methodological issues arising from the study of complex decisions including differences between field versus laboratory research.

Dynamic decision making (DDM) is defined by three common features: a series of actions must be taken over time to achieve some overall goal, the actions are interdependent so that later decisions depend on earlier actions, and the environment changes both spontaneously and as a consequence of earlier actions (Edwards, 1962). Dynamic decision tasks differ from sequential decision tasks (see Diederich, 2001) in that the former are primarily concerned with controlling dynamic systems over time, whereas the latter are more concerned with sequential search for information to be used in making decisions.

Psychological research on DDM began with Toda's (1962) pioneering study of human performance on a game called the 'fungus eater,' in which human subjects controlled a robot's search for uranium and fuel on a hypothetical planet. Subsequently, human performance has been examined across a wide variety of dynamic decision tasks including computer games designed to simulate stock purchases (Ebert, 1972; Rapoport, 1966), welfare management (Dorner, 1980; Mackinnon and Wearing, 1980), vehicle navigation (Jagacinski and Miller, 1978; Anzai, 1984), health management (Kleinmuntz and Thomas, 1987; Kerstholt, 1994), production and inventory control (Serman, 1989; Berry and Broadbent, 1988), supervisory control (Kirluk et al., 1993a,b), firefighting (Brehmer and Allard, 1991; Brehmer and Nahlinger, 2007), prey and predator problems (Jensena and Brehmer, 2003), plumbing and cash flow tasks (Gonzalez and Dutt, 2011; Sweeney and Serman, 2000), and farming on Mars (Gureckis and Love, 2009a). Cumulative progress in this field has been summarized in a series of empirical reviews by Edwards (1962), Rapoport (1975), Funke (1991), Brehmer (1992), Serman (1994), Kerstholt and Raaijmakers (1997), and Jagacinski and Flack (2003).

Stochastic Optimal Control Theory

To illustrate how psychologists study human performance on dynamic decision tasks, consider the following experiment by

You (1989). Subjects were initially presented a 'cover' story describing the task: "Imagine that you are being trained as a psychiatrist, and your job is to treat patients using a psychoactive drug to maintain their health at some ideal level." Subjects were instructed to choose the drug level for each day of a simulated patient after viewing all of a patient's previous records (treatments and health states). Subjects were trained on 20 simulated patients, with 14 days per patient, all controlled by a computer simulation program.

There are a few general points to make about this type of task. First, laboratory tasks such as this are oversimplifications of real-life tasks, designed for experimental control and theoretical tractability. However, more complex simulations also have been studied to provide greater realism (e.g., Brehmer and Allard's, 1991, firefighting task). Second, the above task is an example of a discrete time task (only the sequence of events is important), but real-time simulations have also been examined where the timing of decisions becomes critical (e.g., Brehmer and Allard's, 1991, firefighting task). Third, the cover story (e.g., health management) provides important prior knowledge for solving the task, and so the findings depend both on the abstract task properties as well as the concrete task details (Kleider, 1975). Fourth, the stimulus events are no longer under complete control of the experimenter, but instead they are also influenced by the subject's own behavior. Thus experimenters need to switch from a stimulus-response toward a cybernetic paradigm for designing research (cf. Brehmer, 1992; Rapoport, 1975).

This health management task can be formalized by defining $H(t)$ as the state of the patient's health on day t ; $T(t)$ is the drug treatment presented on day t , and $w(t)$ is a random shock that may disturb the patient on any given day. Figure 1 is a feedback diagram that illustrates this dynamic decision task. In this figure, S represents the environmental system that takes both the disturbance, w , and the decision maker's control action, T , as inputs, and produces the patient's state of health, H , as output. D represents the decision maker's policy that takes both the observed, H , and desired, H^* , states of health as inputs, and produces the control action, T , as output.




Figure 1 A feedback diagram of an example dynamic decision task.

Based on these definitions, this task can be analyzed as a stochastic linear optimal control problem (Rouse, 1980): determine treatments $T(1), \dots, T(N)$, for $N = 14$ days, which minimize the objective function

$$F = E \left\{ \sum_{t=1, N} \alpha [H(t) - H^*(t)]^2 + \beta T^2(t) \right\} \quad [1]$$

contingent upon the linear stochastic dynamic system,

$$H(t+1) = a_1 H(t) + a_2 H(t-1) + b_1 T(t) + b_2 T(t-1) + b_3 T(t-2) + w(t) \quad [2]$$

Standard dynamic programming methods (Bertsekas, 1976) may be used to find the optimal solution to this problem. For the special case where the desired state of health is neutral ($H^* = 0$), some treatment effect takes place the very next day (b_1 is nonzero), and there is no cost associated with the treatments (i.e., $\beta = 0$), the optimal policy

$$T(t) = -[(a_1/b_1)H(t) + (a_2/b_1)H(t-1) + (b_2/b_1)T(t-1) + (b_3/b_1)T(t-2)] \quad [3]$$

is the treatment that forces the mean health state to equal the ideal (zero) on the next day. If the cost of treatment is nonzero (i.e., $\beta > 0$), then the solution is a more complex linear function of the previous health states and past treatments (see Bertsekas, 1976; You, 1989).

Dynamic programming is a general-purpose method that can be used to solve for optimal solutions to many dynamic decision tasks. Although the example above employed a linear control task, dynamic programming can also be used to solve many nonlinear control problems (see Bertsekas, 1976). However, for highly complex tasks, dynamic programming may not be practical, and heuristic search methods such as genetic algorithms (Holland, 1994) may be more useful.

The formal task analysis presented above provides a basis for determining factors that may affect human performance on the task (cf., Brehmer, 1992). One factor is the stability of the dynamic system, which for eqn [2] depends on the two coefficients, a_1 and a_2 . In particular, this system is stable if the roots of the characteristic equation

$$\lambda^2 - a_1 \lambda - a_2 = 0$$

are less than one in magnitude (see Luenberger, 1979). A second factor is the controllability of the system, which depends on the three coefficients, b_1 , b_2 , and b_3 (see Luenberger, 1979). For example, if the treatment effect is delayed ($b_1 = 0$), then the simple control policy shown in eqn [3] is no longer feasible, and the optimal policy is more a complex linear function of the system coefficients.

Consistent with previous research, You (1989) found that even after extensive experience with the task, subjects frequently lost control of their patients and the average performance of human subjects fell far below optimal performance. However, this is a gross understatement. Sterman (1989) found that when subjects tried to manage a simulated production task, they produced costs 10 times greater than optimal, and their decisions induced costly cycles even though the consumer demand was constant. Brehmer and Allard (1991) found that when subjects tried to fight simulated forest fires, they frequently allowed their headquarters to burn down despite desperate efforts to put the fire out. Kleinmuntz and Thomas (1987) found that when subjects tried to manage their simulated patients' health, they often let their patients die while wasting time waiting for the results of nondiagnostic tests and performed more poorly than a random benchmark.

Alternative Explanations for Human Performance

There are many alternative reasons for this suboptimal performance. By their very nature, dynamic decision tasks entail the coordination of many tightly interrelated psychological processes including causal learning, planning, problem solving, and decision making (cf. Toda, 1962). Six different psychological approaches for understanding human DDM behavior have been proposed, each focusing on one of the component processes.

The first approach was proposed by Rapoport (1975), who suggested that suboptimal performance could be derived from an optimal model either by adding information processing constraints on the planning process, or by including subjective utilities into the objective function of eqn [1]. For example, Rapoport (1966, 1967) found that human performance in his stock purchasing tasks was accurately reproduced by assuming that subjects could only plan a few steps ahead (about three steps), as compared to the optimal model with an unlimited planning horizon. In this case, dynamic programming was useful for providing insights about the constraints on human planning capabilities. As another example, Rapoport et al. (1970) attempted to predict performance in a multistage investment game by assuming that the utility function was a concave function of monetary value. In this case, Rapoport et al. (1970) used dynamic programming to derive an elegant decision rule, which predicted that investments should be independent of the size of the capital accumulated during play. Contrary to this prediction, human decision makers were strongly influenced by this factor, and so in this case, dynamic programming was useful for revealing empirical flaws with the theory.

An alternative approach, proposed by Brehmer (1992) and Sterman (1994), is that suboptimal performance is caused by a misconception of the dynamic system (eqn [2]). In other words, a subject's internal model of the system does not match the true model. In particular, human performers seem to have great difficulty discerning the influence of delayed feedback and understanding the effects of nonlinear terms in the system (Rahmandad et al., 2009). Essentially, subjects solve the problem as if it was a linear system that has only zero lag terms. In the case of eqn [2], the subjective decision policy is simply

$$T(t) = -c_1 H(t)$$

where c_1 is estimated from a subject's control decisions. [Sterman \(1989\)](#) and [Diehl and Sterman \(1993\)](#) found that this type of simplified subjective policy described their subjects' behavior very accurately.

Several recent studies have investigated how to improve people's understanding of the dynamic systems in DDM tasks. [Kopainsky and Sawicka \(2011\)](#) used an interactive simulation tool, which allowed participants to explore the dynamic features of a resource management task. They found that the tool allowed decision makers to construct more accurate mental representations of the task, resulting in greater understanding and improved performance. [Cronin, Gonzalez, and Sterman \(2008\)](#) suggested that having participants verbalize their decisions enhanced their performance in an accumulator task. [Sterman \(2010\)](#) reported that training MIT graduate students in system dynamics (e.g., with an introductory system dynamics course) also improved their overall task performance.

A more general method for estimating subjective decision policies was proposed by [Jagacinski \(Jagacinski and Miller, 1978; Jagacinski and Hah, 1988\)](#). Consider once again the example problem employed by [You \(1989\)](#). In this case, a subject's treatment decision on each trial $T(t)$ could be represented by a linear control model:

$$T(t) = c_0 + c_1H(t) + c_2H(t-1) + c_3H(t-2) + c_4T(t) + c_5T(t-1) + c_6T(t-2) + \text{error}$$

where the subjective coefficients (c_0, c_1, \dots, c_6) are estimated by a multiple regression analysis. This is virtually the same as performing a 'lens model' analysis to reveal the decision maker's policy (see [Kleinmuntz, 1993; Slovic and Lichtenstein, 1971](#)). This approach has been successfully applied in numerous applications ([Jagacinski and Miller, 1978; Jagacinski and Hah, 1988; Kerlik, Miller, Jagacinski, 1993; Kleinman et al., 1980; You, 1989](#)). Indeed, [You \(1989\)](#) found that subjects made use of both lags 1 and 2 for making their treatment decisions.

Heuristic approaches to strategy selection in dynamic decision tasks have been explored by [Kleinmuntz \(1985\)](#), [Kleinmuntz and Thomas \(1987\)](#), and [Kerstholt \(1994, 1996\)](#). These researchers examined health management tasks that entailed dividing time between two strategies: collecting information from diagnostic tests before choosing a treatment, versus treating patients immediately without conducting any more diagnostic tests. A general finding is that even experienced subjects tend to overuse information collection, resulting in poorer performance than that which could be obtained from a pure treatment (no test) strategy. This finding seems to run counter to the adaptive decision-making hypothesis of [Payne et al. \(1993\)](#), which claims that subjects prefer to minimize effort and maximize performance. The information collection strategy is both more effortful and less effective than the treatment strategy in this situation. In another study, [Cronin et al. \(2008\)](#) employed simple stock-flow problems and found that people's inability to understand the dynamics of the task resulted from their use of inappropriate heuristics, which sometimes yielded correct responses, but fundamentally misrepresented the interplay of task variables. Performance remained poor, regardless of the amount of data presented to the participants, the method of presentation, and changes to the cover stories or context.

Finally, an individual difference approach to understanding performance on complex dynamic decision tasks was developed by [Dorner and his colleagues \(see Funke, 1991 for a review\)](#). Subjects are divided into two groups (good versus poor) on the basis of their performance on a complex dynamic decision task. Subsequently, these groups are compared on various behaviors to identify the critical determinants of performance. This research indicates that subjects who perform best are those that set integrative goals, collect systematic information, and evaluate progress toward these goals. Subjects who tend to shift from one specific goal to another, or focus exclusively on only one specific goal, perform more poorly. Along these lines, [Bisantz et al. \(2000\)](#) applied the lens model to DDM tasks and found that much of the differences between the best and worst performers were due to the relative consistency with which top performers made judgments, and not to different judgment policies, nor to characteristics of the judgment environments.

Learning to Control Dynamic Systems

Although human performers remain suboptimal even after extensive task training, almost all past studies reveal systematic learning effects. First, overall performance often rapidly improves with training (see e.g., [Brehmer, 1992; Gonzalez and Quesada, 2003; Mackinnon and Wearing, 1985; Rapoport, 1966; Sterman, 1989](#)). Furthermore, subjective policies tend to evolve over trial blocks toward the optimal policy ([Jagacinski and Miller, 1978; Jagacinski and Hah, 1988; You, 1989](#)). Therefore, learning processes are important for explaining much of the variance in human performance on dynamic decision tasks (cf., [Hagmayer et al., 2010; Hogarth, 1981](#)). Three different frameworks for modeling human learning processes in dynamic decision tasks have been proposed ([Osman, 2010](#)).

A production rule model was developed by [Anzai \(1984\)](#) to describe how humans learn to navigate a simulated ship. The general idea is that past and current states of the problem are stored in working memory. Rules for transformation of the physical and mental states are represented as condition-action type production rules. A product rule fires whenever the current state of working memory matches the conditions for a rule. When a production rule fires, it deposits new inferences or information in memory, and it may also produce physical changes in the state of the system. Navigation is achieved by a series of such recognize-act cycles of the production system. Learning occurs by creating new production rules that prevent earlier erroneous actions, and new rules are formed on the basis of means-ends type of analyses for generating subgoal strategies. Simulation results indicated that the model learned strategies similar to those produced in the verbal protocols of subjects, although evaluation of the model was dependent on qualitative judgments rather than quantitative measurements.

An instance- or exemplar-based model was later developed by [Dienes and Fahey \(1995\)](#) to describe how humans learn to control a simulated sugar production task, and to describe how they learn to manage the emotional state of a hypothetical person. This model assumes that whenever an action leads to a successful outcome, the preceding situation and the successful

response are stored together in memory. On any given trial, stored instances are retrieved on the basis of similarity to the current situation, and the associated response is applied to the current situation. This model was compared to a simple rule-based model like that employed by Anzai (1984). The results indicated that the exemplar learning model produced more accurate predictions for delayed feedback systems, but the rule-based model performed better when no feedback delays were involved. This conclusion agrees with earlier ideas presented by Berry and Broadbent (1988) that delayed feedback tasks involve implicit learning process, whereas tasks without delay are based on explicit learning processes.

Gonzalez et al. (2003) proposed the more sophisticated instance-based learning theory (IBLT), which posits that decision makers adapt their judgment strategies from heuristic-based to instance-based as they accumulate and refine their knowledge and learn to recognize situations based on similarity to past instances. This model was used to explain data from a water distribution simulation. In the task, participants managed a water purification plant by shifting water between several tanks. As rain fell some tanks spontaneously filled with water and had to be pumped. IBLT construes the decision maker as learning to control this DDM task through the accumulation and refinement of instance-based knowledge. Participants begin by operating according to simple, domain-general heuristics. However, as decision makers gain experience in the task, they accumulate instance knowledge that can be used to improve performance.

Figure 2 gives a graphical representation of how the model performs a DDM task. The process begins with a search through memory for instances that match the current state. Instances are represented as situation–decision–utility (SDU) triplets, which code environmental cues, the decision made in that environment, and the utility of the experienced outcome. SDUs are retrieved based on the similarity between the current state and each instance in memory. If no sufficiently similar instances are found, a heuristic is used to choose an action. If a similar entry is found in memory, the model must decide whether to search for more alternatives or to execute the current best action. Additional alternatives are evaluated until one exceeds the decision maker's aspiration threshold or time runs out before a decision deadline, at which point the action with the highest

utility is chosen. After executing an action, memory is updated. New SDUs are accumulated as more alternatives and more decision situations are confronted. Over time, the model is able to use more and more diagnostic cues to search its growing memory. Feedback on previous actions is used to update the utilities of the SDUs. With experience, the model transitions away from heuristic-based decisions as it adapts to the dynamic environment. Gonzalez et al. (2003) tested several versions of IBLT and found that the full version captured participants' changes in performance as they moved away from random or heuristic-based decisions, toward instance-based strategies. Gonzalez and Dutt (2011) used IBLT to model subjects' behavior in two different versions of a decision-from-experience task. Sampling and repeated-choice paradigms are often thought to rely on differences in psychological mechanisms, but IBLT was able to account for choices in both scenarios using a common set of cognitive mechanisms.

Gibson et al. (1997) used an artificial neural network model to describe learning in a sugar production task. This model receives input representing the current state of the environment and the current goal for the task. Hidden layers compute the next action and predict its consequences. Connections between hidden layers and input and output layers are learned by back propagating prediction errors and deviations between observed outcomes and the goal state. This model provided good accounts of subjects' performance during both training and subsequent generalization tests under novel conditions. Recently, Gibson (2007) extended this model by assuming that decision makers first sort their evidence into categories before weighting and using the sorted evidence. This two-stage model better predicted decision makers' performance in a new, related bargaining task. Unfortunately, no direct comparisons with rule- or exemplar-based learning models were conducted, and this remains a challenge for future research.

Another approach to modeling the learning of dynamic systems is the reinforcement learning model (see Sutton and Barto, 1998, for a review of this approach). The goal of this algorithm is to maximize the future cumulative reward through interacting with the environment. Consequently it does not need to know the model of the environment or the correct input/output pairs (as in supervised learning). Optimization is based on the agent's online samples from the environment.

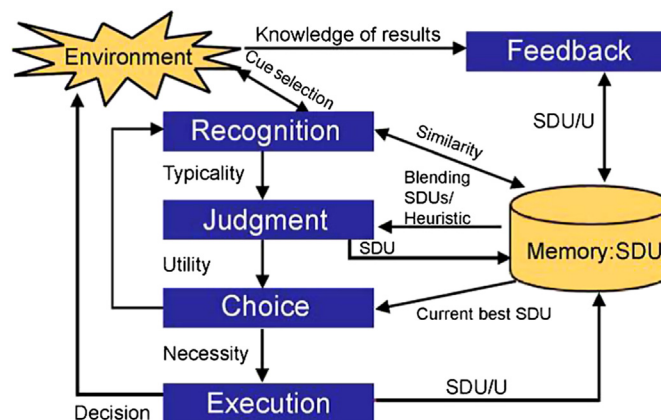


Figure 2 IBLT model diagram.

Recently, Gureckis and Love (2009b) used model-based reinforcement learning to examine human behavior in a 'Farming on Mars' task. Participants' goal was to use two Martian robots to establish oxygen for later human residents. Only one robot could be used at a time, and each performed differently. Participants had to learn about each robot's oxygen output through experience. When rewards were generated probabilistically based on subjects' choices, Gureckis and Love found that decision makers preferred the robot with high short-term output but low long-term output. In the continuous reward condition, where rewards were a deterministic function of subjects' choices, decision makers preferred the robot with the better long-term output but worse short-term output.

In another study, Gureckis and Love (2009a) examined noise effects by implementing noise on both rewards and representations of the current state of task. They found that some moderate level of reward noise caused participants to explore different strategies and helped them find the optimal policy among their alternatives. However, noisy state cues decreased performance relative to no noise conditions. It seems that decision makers use state information as support for their decisions, and so noisy state information is more harmful than noisy reward. Continuing with the 'Farming on Mars' task, Otto and Love (2011) used a reinforcement learning framework to model the impacts of forgone rewards on DDM. They found that providing more information about what could have been gained at each point interfered with subjects' capability to make long-term choices.

Laboratory versus Naturalistic Decision Research

Klein and his associates (Klein et al., 1993; Klein, 2008; Lipshitz et al., 2001; Zsombok and Klein, 1997) have made progress toward understanding dynamic decisions in applied field research settings, which they call naturalistic decision making (e.g., interview fire chiefs after fighting a fire). Field research complements laboratory research in two ways: on the one hand, it provides a reality check on the practical importance of theory developed in the laboratory; on the other hand, it provides new insights that can be tested with more control in the laboratory.

The general findings drawn from naturalistic decision research provide converging support for the general theoretical conclusions obtained from the laboratory. First, it is somewhat of a misnomer to label this kind of research 'decision making' because decision processes comprise only one of the many cognitive processes engaged by these tasks – learning, planning, and problem solving are just as important. Decision making is used to define the overall goal, but then the sequence of actions follows a plan that has either been learned in the past or generated by a problem-solving process. Second, learning processes may explain much of the variance in human performance on dynamic decision tasks. Decision makers use the current goal and current state of the environment to retrieve actions that have worked under similar circumstances in the past. Klein's (1998) recognition-primed decision model is based on this principle, and this same basic idea underlies the production rule, exemplar, and neural network learning models. Third, learning from extensive experience is the key to

understanding why novel subjects fail where experts succeed. A few hundred trials in a laboratory task are relatively insignificant in comparison with say 25 years of experience on a job. Naïve subjects may fail to prevent their headquarters from burning down in a simulated forest fire, but expert fire chiefs succeed in saving our national parks from forest fires every year.

See also: Decision and Choice: Heuristics; Decision and Choice: Sequential Decision Making; Learning: Mathematical Learning Theory, History of; Mathematical Psychology; Optimal Control Theory; Problem Solving and Reasoning, Psychology of.

Bibliography

- Anzai, Y., 1984. Cognitive control of real-time event driven systems. *Cognitive Science* 8, 221–254.
- Berry, D.C., Broadbent, D.E., 1988. Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology* 79, 251–272.
- Bertsekas, D.P., 1976. *Dynamic Programming and Stochastic Control*. Academic Press, N.Y.
- Bisantz, A.M., Kirlik, A., Gay, P., Phipps, D.A., Walker, N., Fisk, A.D., 2000. Modeling and analysis of a dynamic judgment task using a Lens model approach. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans* 6, 605–616.
- Brehmer, B., 1992. Dynamic decision making: human control of complex systems. *Acta Psychologica* 81, 211–241.
- Brehmer, B., Allard, R., 1991. Real-time dynamic decision making: effects of task complexity and feedback delays. In: Rasmussen, J., Brehmer, B., Leplat, J. (Eds.), *Distributed Decision Making; Cognitive Models for Cooperative Work*. Wiley, Chichester.
- Brehmer, B., Nahlinder, S., 2007. Achieving what cannot be done: coping with the time constants in a dynamic decision task by doing something else. *Scandinavian Journal of Psychology* 48, 359–365.
- Cronin, M., Gonzalez, C., Sterman, J.D., 2008. Why don't well-educated adults understand accumulation? A challenge to researchers, educators, and citizens. *Organizational Behavior and Human Decision Processes* 101, 116–130.
- Diederich, A., 2001. *Decision and Choice: Sequential Decision Making*. International Encyclopedia of the Social and Behavioral Sciences: Methodology, Mathematics, and Computer Science. Pergamon, Amsterdam.
- Diehl, E., Sterman, J.D., 1993. Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes* 62, 198–215.
- Dienes, Z., Fahey, R., 1995. Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 21, 848–862.
- Dorner, D., 1980. On the problems people have in dealing with complexity. *Simulation and Games* 11, 87–106.
- Ebert, R.J., 1972. Human control of a two-variable decision system. *Organizational Behavior and Human Performance* 7, 237–264.
- Edwards, W., 1962. Dynamic decision theory and probabilistic information processing. *Human Factors* 4, 59–73.
- Funke, J., 1991. Solving complex problems: exploration and control of complex systems. In: Sternberg, R.J., Frensch, P.A. (Eds.), *Complex Problem Solving: Principles and Mechanisms*. Erlbaum, Hillsdale, NJ.
- Gibson, F., Fichman, M., Plaut, D.C., 1997. Learning in dynamic decision tasks: computational model and empirical evidence. *Organizational Behavior and Human Performance* 71, 1–35.
- Gibson, F., 2007. Learning and transfer in dynamic decision environments. *Computational and Mathematical Organization Theory* 13, 39–61.
- Gonzalez, C., Dutt, V., 2011. A generic dynamic control task for behavioral research and education. *Computers in Human Behavior* 27, 1904–1914.
- Gonzalez, C., Lerch, F.J., Lebiere, C., 2003. Instance-based learning in dynamic decision making. *Cognitive Science* 27, 591–635.
- Gonzalez, C., Quesada, J., 2003. Learning in dynamic decision making: the recognition process. *Computational and Mathematical Organization Theory* 9, 287–304.

- Gureckis, T.M., Love, B.C., 2009a. Learning in noise: dynamic decision making in a variable environment. *Journal of Mathematical Psychology* 53, 180–193.
- Gureckis, T.M., Love, B.C., 2009b. Short-term gains, long-term pains: how cues about state aid learning in dynamic environments. *Cognition* 113, 293–313.
- Hagmayer, Y., Meder, B., Osman, M., Mangold, S., Lagnado, D., 2010. Spontaneous causal learning while controlling a dynamic system. *The Open Psychology Journal* 3, 145–162.
- Hogarth, R.M., 1981. Beyond discrete biases: functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin* 90, 197–217.
- Holland, J., 1994. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge.
- Jagacinski, R.J., Flach, J.M., 2003. *Control Theory for Humans: Quantitative Approaches to Modeling Performance*. Erlbaum, Mahwah, NJ.
- Jagacinski, R.J., Hah, S., 1988. Progression-regression effects in tracking repeated patterns. *Journal of Experimental Psychology: Human Perception and Performance* 14, 77–88.
- Jagacinski, R.J., Miller, R.A., 1978. Describing the human operator's internal model of a dynamic system. *Human Factors* 20, 425–433.
- Jensena, E., Brehmer, B., 2003. Understanding and control of a simple dynamic system. *System Dynamic Review* 19, 119–137.
- Kerstholt, J.H., 1994. The effect of time pressure on decision making behavior in a dynamic task environment. *Acta Psychologica* 86, 89–104.
- Kerstholt, J.H., 1996. The effects of information costs on strategy selection in dynamic tasks. *Acta Psychologica* 94, 273–290.
- Kerstholt, J.H., Raaijmakers, J.G.W., 1997. Decision making in dynamic task environments. In: Crozier, W.R., Svenson, O. (Eds.), *Decision Making: Cognitive Models and Explanations*. Routledge, London, pp. 205–217.
- Kirlik, A., Miller, R.A., Jagacinski, R.J., 1993a. Supervisory control in a dynamic and uncertain environment: process model of skilled human-environment interaction. *IEEE Transactions on Systems, Man, and Cybernetics* 23, 929–952.
- Kirlik, A., Plamondon, D.D., Lytton, L., Jagacinski, R.J., 1993b. Supervisory control in a dynamic and uncertain environment: laboratory task and crew performance. *IEEE Transactions on Systems, Man, and Cybernetics* 23, 1130–1138.
- Klein, G., 1998. *Sources of Power: How People Make Decisions*. MIT Press, Cambridge.
- Klein, G., Orasanu, J., Calderwood, R., Zsombok, C.E., 1993. *Decision Making in Action: Models and Methods*. Ablex, Norwood, NJ.
- Klein, G., 2008. Naturalistic decision making. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 50, 456–460.
- Kleinman, D.L., Pattipati, K.R., Ephrath, A.R., 1980. Quantifying an internal model of target motion in a manual tracking task. *IEEE Transactions on Systems, Man, and Cybernetics* 10, 624–636.
- Kleinmuntz, D., 1985. Cognitive heuristics and feedback in a dynamic decision environment. *Management Science* 31, 680–702.
- Kleinmuntz, D., 1993. Information processing and misperceptions of the implications of feedback in dynamic decision making. *System Dynamics Review* 9, 223–237.
- Kleinmuntz, D., Thomas, J., 1987. The value of action and inference in dynamic decision making. *Organizational Behavior and Human Decision Processes* 39, 341–364.
- Kleiter, G.D., 1975. Dynamic decision behavior: comments on Rapoport's paper. In: Wendt, D., Vlek, C. (Eds.), *Utility, Probability, and Human Decision Making*. Reidel, Dordrecht-Holland, pp. 371–380.
- Kopainsky, B., Sawicka, A., 2011. Simulator-supported descriptions of complex dynamic problems: experimental results on task performance and system understanding. *System Dynamics Review* 27, 142–172.
- Lipshitz, R., Klein, G., Orasanu, J., Salas, E., 2001. Taking stock of naturalistic decision making. *Journal of Behavioral Decision Making* 14, 332–351.
- Luenberger, D.G., 1979. *Introduction to Dynamic Systems*. Wiley, NY.
- Mackinnon, A.J., Wearing, A.J., 1980. Complexity and decision making. *Behavioral Science* 25, 285–296.
- Mackinnon, A.J., Wearing, A.J., 1985. Systems analysis and dynamic decision making. *Acta Psychologica* 58, 159–172.
- Osman, M., 2010. Controlling uncertainty: a review of human behavior in complex dynamic environments. *Psychological Bulletin* 136, 65–86.
- Otto, A.R., Love, B.C., 2011. You don't want to know what you're missing: when information about forgone rewards impedes dynamic decision making. *Judgment and Decision Making* 5, 1–10.
- Payne, J.W., Bettman, J.R., Johnson, E.J., 1993. *The Adaptive Decision Maker*. Cambridge University Press, NY.
- Rahmandad, H., Repenning, N., Sterman, J., 2009. Effects of feedback delay on learning. *System Dynamics Review* 25, 309–338.
- Rapoport, A., 1966. A study of human control in a stochastic multistage decision task. *Behavioral Science* 11, 18–32.
- Rapoport, A., 1967. Dynamic programming models for multistage decision making. *Journal of Mathematical Psychology* 4, 48–71.
- Rapoport, A., 1975. Research paradigms for studying dynamic decision behavior. In: Wendt, D., Vlek, C. (Eds.), *Utility, Probability, and Human Decision Making*. Reidel, Dordrecht-Holland, pp. 347–369.
- Rapoport, A., Jones, L.V., Kahan, J.P., 1970. Gambling behavior in multiple-choice multi-stage betting games. *Journal of Mathematical Psychology* 7, 12–36.
- Rouse, W.B., 1980. *Systems Engineering Models of Human-Machine Interaction*. North-Holland, NY.
- Slovic, P., Lichtenstein, S., 1971. Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance* 6, 649–744.
- Sterman, J.D., 1989. Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes* 43, 301–335.
- Sterman, J.D., 1994. Learning in and about complex systems. *System Dynamics Review* 10, 291–330.
- Sterman, J.D., 2010. Does formal system dynamics training improve people's understanding of accumulation? *System Dynamics Review* 26, 316–334.
- Sutton, R.S., Barto, A.G., 1998. *Reinforcement Learning*. MIT Press, Cambridge, MA.
- Sweeney, L.B., Sterman, J.D., 2000. Bathtub dynamics: initial results of a systems thinking inventory. *System Dynamic Review* 16, 249–286.
- Toda, M., 1962. The design of the fungus eater: a model of human behavior in an unsophisticated environment. *Behavioral Science* 7, 164–183.
- You, 1989. *Disclosing the decision-maker's internal model and control policy in a dynamic decision task using a system control paradigm*. Unpublished MA thesis, Purdue University.
- Zsombok, C.E., Klein, G., 1997. *Naturalistic Decision Making*. Erlbaum, Mahwah, NJ.