Oxford Handbooks Online

Decision-Making: A Cognitive Science Perspective

Cleotilde Gonzalez

The Oxford Handbook of Cognitive Science
Edited by Susan E. F. Chipman

Print Publication Date: Oct 2017 Subject: Psychology, Cognitive Psychology
Online Publication Date: Nov 2014 DOI: 10.1093/oxfordhb/9780199842193.013.6

Abstract and Keywords

This chapter overviews topics in judgment and decision making from a cognitive science perspective. It advocates a "closed-loop" view of decision making: an interactive and continuous dynamic process of exchanges between humans and their environment. The chapter first discusses the "open-loop" view of decision making that has dominated the field for many decades, beginning with a historical perspective on rationality and bounded rationality to distinguish the closed and open-loop views and the research from two major fields that study decision making: economics and psychology. It then presents foundational research for the closed-loop view that involves probability learning and dynamic decision making, adaptive decision making, and recent research on dynamic decision making and decisions from experience. The last section presents the naturalistic decision-making perspective and its connections to cognitive engineering and human factors. It concludes with a view on future research at individual, team, group, and societal levels.

Keywords: decision making, cognitive science, dynamic decision making, learning, adaptive decision making, naturalistic decision making

Decision making is a "high-level" cognitive process that is clearly distinguishable from other processes in at least two ways: it builds on more basic cognitive processes such as perception, memory, and attention, and it is uniquely identified by its essential element: the process of *choice*. Choice is the act of selecting among alternatives, whether they are present at the same time or they develop over time. The process of choice is highly influenced by the cognitive processes that occur before a choice is made (e.g., perception, recognition, and judgment) and those that occur after a choice is made (e.g., feedback and learning). Notably, the *judgment* process, which precedes choice, involves evaluating the merits of and preferences for different alternatives. These two processes, judgment and choice, have been the main focus of the study of decision making, a field often referred to as "judgment and decision making." As it will become apparent throughout this chapter, judgment and choice present only a partial view of all the cognitive processes involved in decision making, and most of the other processes, such as perception, recognition, feedback, and learning, have been relatively neglected in the literature of decision sciences.

The cognitive science perspective on decision making that I advocate is an interactive view, one highly influenced by roots in control engineering and ecological psychology (Flach, Hancock, Caird, & Vincente, 1995; Pew & Baron, 1978; Wickens & Kramer, 1985). Under this view, a human (the decision maker) collaborates with an environment in order to accomplish a task through repeated decisions. A decision maker perceives information from the environment and transforms that information to find and create alternatives, build preferences, and evaluate options that lead to a choice. An action is executed and it naturally results in changes in the environment. Then, feedback (i.e., the knowledge of outcomes from actions taken) must be processed in order to reinforce (p. 250) or not past decisions (i.e., learn from past choices). I refer to this process as a *closed-loop view* of decision making (Gonzalez, 2012), and it is the principal flow of events behind most models of dynamic decision making and models of decisions from experience in dynamic situations, which will be explained below.

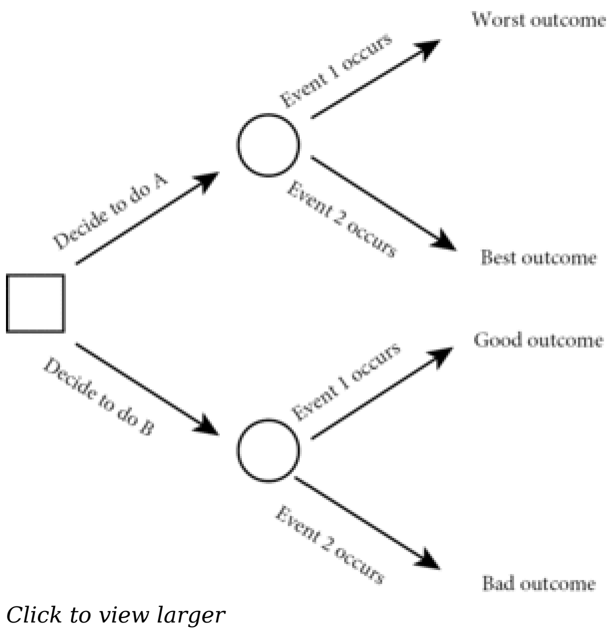

Click to view larger

Fig. 13.1 The open-loop, linear model of decision making.

Although this interactive perspective would seem intuitive and natural for most cognitive psychologists, this has not been the dominant model of information processing in decision sciences. Instead, an *open-loop* linear model, such as that illustrated in Figure 13.1, has been the most common conceptualization of decision making in the past decades (Gonzalez, 2012; Hastie, 2001). This conceptualization involves (1) the explicit presentation of choice options or alternatives, often represented as “branches” of a decision tree; (2) beliefs about objective events in the world, which represent uncertainty in the environment (often described as probabilities or likelihoods); and (3) desires or utilities that represent the consequences associated with the outcomes of each action–event combination or subjective evaluative reactions associated with each outcome.

In what follows, I briefly discuss a historical view of decision-making research that should help readers understand the open-loop view, as well as encourage the closed-loop view that I take in the dynamic decision making perspective. The open-loop view has largely dominated decision-making research; thus, I discuss the most traditional approach—heuristics and biases—before introducing dynamic decision making and cognitive models of decisions from experience in dynamic environments. In the final section, I discuss the naturalistic decision-making perspective and provide a personal view of the future of decision-making research.

Maximizing and Satisficing, Rational and Irrational, Economics and Psychology: The Open-Loop View of Decision Making

Theories that explain human decision making have traditionally involved principles and developments taken from economics and psychology. These two disciplines have proposed what appear to be conflicting mechanisms and explanations for decision making. On the one hand, economists have often assumed humans to be utility maximizers (i.e., "rational"), whereas psychologists have aimed at demonstrating the many different decision situations in which humans do not maximize utility (i.e., "irrational"). Given that most cognitive treatments of decision making depart from the concept of rationality, I discuss some of these ideas next. However, there have been many extensive discussions on this theme. Most are excellent accounts of the history of decision sciences that depart from theories of expected utility (which is considered to include theories of rational behavior) to more recent trends of cognitive decision-making research (e.g., Einhorn & Hogarth, 1981; Goldstein & Hogarth, 1997; Griffin, Gonzalez, Koehler, & Gilovich, 2012). Thus, I only discuss those relevant historical (p. 251) aspects that have contributed to the formation and prevalence of the open-loop view of decision making in order to contrast it to the cognitive science perspective that I emphasize—the closed-loop view.

Rationality is a central concept in the discipline of economics and to the principle of maximization of subjective expected utility (SEU). According to this principle, a decision maker evaluates the attractiveness of an option by combining the probability of each possible outcome and the subjective utility or personal value of each outcome. Since its initial proposal, the SEU theory's value as a principle of human decision making has been heavily criticized. A notable example is the critique offered early on by Herbert Simon

(1955, 1957), in which he described the discrepancies between the SEU model and the reality of human behavior and proposed an "approximate" rationality or *satisficing* mechanism for decision making (i.e., "bounded rationality"). Under the bounded rationality mechanism, people can adapt to their environment by identifying actions that are only satisfactory to their goals and that could be applied without applying cognitively demanding and sophisticated rules that humans are unable to use.

Another major research breakthrough was gained from human behavioral experiments that questioned the fundamental assumptions of the SEU model. The work of Kahneman, Tversky, and their colleagues helped to shift attention from examples that merely dispute SEU theory to providing explanations of *how* people make decisions, as described by *prospect theory* (Kahneman & Tversky, 1979). This theory has been a prominent model used to explain and generalize deviations from expected utility theory.

When demonstrating the explanatory power of prospect theory, researchers have traditionally used monetary gambles (i.e., "prospects") that explicitly state outcomes and associated probabilities. People are presented with a description of the alternatives, and they are asked to make a choice based on the conditions described. They are asked to make *decisions from description*. For example:

> Which of the following would you prefer?
>
> **A:** a .8 chance to get $4 and .2 chance to get $0
>
> **B:** get $3 for sure

Using descriptive prospects, researchers have investigated a large number of situations in which people behave against utility maximization and in agreement with prospect theory, producing an impressive list of "heuristics and biases" (Kahneman, Slovic, & Tversky, 1982; Tversky & Kahneman, 1974). Through the years, these consistent deviations from rational behavior have been identified, replicated, and extended using laboratory experiments, to the extent that this type of research has dominated the field for the past six decades.

Heuristics and Biases

The strength of the heuristics and biases approach is that each of the demonstrations of "irrationality" has a clear baseline for comparison (a normative model or a description of optimal human behavior, being either utility maximization or optimality in other forms, such as expert behavior). Heuristics are the "shortcuts" that humans use to reduce task complexity in judgment and choice, and biases are the resulting gaps between normative behavior and the heuristically determined behavior (Kahneman et al., 1982).

The list of heuristics and biases keeps growing. A summary of the cognitive biases that arise from reliance on judgment heuristics is presented in Table 13.1. This is a partial list of well-known heuristics and biases described in Kahneman et al.'s (1982) book, that highlight three heuristics that are employed in making judgments under uncertainty: representativeness, availability, and anchoring and adjustment. These three heuristics result in at least 13 biases or errors in situations of uncertainty.

Using a similar and related set of demonstrations, Gigerenzer and Todd (1999) presented human use of heuristics as inference mechanisms that can be simple and successful to the degree that they are ecologically rational (adapted to the environment in which decisions are made). Thus, rather than conceptualizing heuristics as the source of "errors" or biases, Gigerenzer and Todd explored the world of human and environment adaptation and argued for the "rationality" of these heuristics. This research program has been quite successful, particularly in highlighting the value of adaptation to the environment. Their concept of "ecological rationality" aims at understanding environmental structures and how heuristics may succeed or fail in particular situations. For an updated view of this perspective, see Gigerenzer, Hertwig, and Pachur (2011).

Information Processing and Process-Tracing Methods

Despite many years of effort investigating heuristics and biases in decision making, we still have only limited answers to the question of *how* people (p. 252) actually go about making decisions. Rather, most research under these programs has been aimed at demonstrating how people *don't* make decisions. The large collection of cognitive biases cannot all be explained by one comprehensive theory, and, most importantly, we do not know how the biases develop and how they emerge in the first place. As a result, little is known of how to prevent them. Most empirical studies to date focus on observable processes, such as choice selection, and ignore cognitive processes that lead to choice, such as recognizing alternatives, deciding when to search for information, evaluating and integrating possible outcomes, and learning from good and bad decisions.

Table 13.1 A Summary of Cognitive Heuristics and Biases

Heuristic	**Biases**
Representativeness: probabilities are assessed by the degree of representativeness (similarity)	**1.** Insensitivity to prior probability of outcomes **2.** Insensitivity to sample size **3.** Misconceptions of chance **4.** Insensitivity to predictability **5.** The illusion of validity **6.** Misconceptions of regression
Availability: probabilities are assessed by the ease with which instances or occurrences are brought to mind	**1.** Biases due to the retrievability of instances **2.** Biases due to the effectiveness of a search set **3.** Biases of imaginability **4.** Illusory correlation
Adjustment and Anchoring: estimates are made by starting from an initial value that is adjusted to yield the final answer.	**1.** Insufficient adjustment **2.** Biases in the evaluation of conjunctive and disjunctive events **3.** Anchoring in the assessment of subjective probability distributions

In decision sciences, attention to the processes that precede choice has been rare but productive. Most of the work in this area has been devoted to the methods involved in the study of the decision process, relying on process-tracing measures and think-aloud protocols (Crozier & Ranyard, 1997). A notable example is mouse-lab studies that emerged from adaptive decision maker studies (Payne, Bentmann, & Johnson, 1993). With a process-tracing tool, “mouse-lab,” researchers monitored the information acquisition process by hiding information behind boxes and requesting that a decision maker explicitly select a box to reveal the information behind it. This method has been successful in providing some insights into the process behind decision making that would otherwise stay eclipsed by a focus on the outcomes of explicitly selected choices. Regardless of the success of these methods, there is still much more to investigate about cognitive processes that precede and follow choice when people make decisions.

The interactive and closed-loop view of decision making is now an influential model in cognitive psychology that addresses what takes place between the presentation of a stimulus and the feedback received from the execution of a choice. This has been illuminated by research on human learning and dynamic decision-making research that started in the 1950s. More recently, research on decisions from experience has been explicitly and successfully contrasted with findings from traditional decisions from description. This perspective is discussed next, under the umbrella of dynamic decision making.

Dynamic Decision Making: A Closed-Loop View of Decision Making

The open-loop view of decision making just described makes a set of assumptions regarding information availability and a decision maker’s ability to process such information. It assumes that the alternatives, presented at the same time, are explicitly described and include the probabilities assigned to each possible outcome. Simon (1955) observed that these conditions would rarely be representative of actual choice situations and that human abilities may not meet the processing demands. The heuristics and biases approach is incomplete and does not provide a full account of the decision-making processes. Simon (1955) highlighted the importance of learning also in the choice process, and decision theory in the 1950s only addressed gambles and simple choice involving skill. For example, (p. 253) probability learning was extensively studied in the 1950s and 1960s (Estes, 1964, 1976; Lee, 1971; Shanks, Tunney, & McCarthy, 2002) in paradigms involving the prediction of the occurrence of two mutually exclusive events in which the decision maker receives feedback about which event took place. These learning experiments were the origin of dynamic decision research (Edwards, 1962).

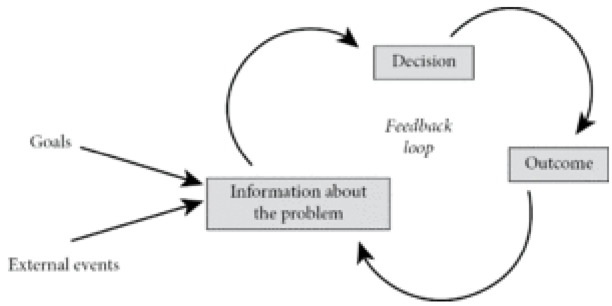

Click to view larger

Fig. 13.2 Closed-loop view of decision making.

Dynamic decision making (DDM) may be conceptualized as a control process: a closed-loop learning process in which decisions are influenced by goals and external events and are the result of previous decisions and previous outcomes. Under this view, decision making is a learning process in which decisions are made based on experience and are feedback-dependent (see Figure 13.2). Alternatives are not presented at the same time, but rather they unfold over time. Decision making is a learning loop: decisions depend on previous choices and also on external events and conditions.

A recent development in decision sciences has expanded on the initial view of choice as a learning process. It has great potential to expand our understanding and provide insights into the dynamic decision making process. The development involves a shift of attention to how decisions are made based on experience (i.e., *decisions from experience*), rather than based on explicit descriptions of options. Researchers use experimental paradigms that involve repeated decisions rather than one-shot decisions, the estimation of possible outcomes and probabilities based on the observed outcomes rather than from a written description, and learning from feedback. All these are natural processes for making decisions in many real-world situations where alternatives, outcomes, and probabilities are unknown. The experimental paradigm often involves two alternatives, represented as two unlabeled buttons, each representing a probability distribution of outcomes unknown to participants. Clicking a button yields an outcome as a result of a random draw from the alternative's distribution. Although there are multiple paradigms for the study of decisions from experience (Gonzalez & Dutt, 2011; Hertwig & Erev, 2009), a common paradigm is the "sampling" paradigm (see Figure 13.3), in which people are able to explore the outcomes of the options without real consequences before they decide to make a final choice.

A key observation that contributed to the initial success of decisions from experience's theoretical development was the "description-experience gap" (Hertwig, Barron, Weber, & Erev, 2004): that the choice an individual makes depends on how information about the problem is acquired (from description or experience), particularly in problems involving outcomes with low probabilities (less than .2; i.e., "rare events"). A robust finding across a range of paradigms for decisions from experience is that people behave as if rare events have *less* impact than they deserve according to their objective probabilities. More importantly, this finding contradicts the prediction from prospect theory that people behave as if rare events have *more* impact than they deserve. However, this theory only applies to "simple prospects with monetary values and stated probabilities" (Kahneman & Tversky, 1979, p. 274). Thus, although prospect theory seems to provide good explanations for decisions from description, findings from experiments about decisions from experience may contradict those predictions in many cases (Hertwig, 2012).

Prospect theory (Kahneman & Tversky, 1979) has been a prominent model to explain human choice behavior in descriptive choices, but a comprehensive model that can explain decisions from experience has (p. 254) not yet been found. In fact, a challenge in understanding the cognitive processes involved in making decisions from experience is the proliferation of highly task-specific models that often predict behavior in a particular task but fail to also explain behavior in other closely related tasks (see discussions in Gonzalez & Dutt, 2011; Lejarraga, Dutt, & Gonzalez, 2012). Gonzalez and colleagues have attempted to address this challenge by providing multiple demonstrations of how cognitive computational models based on one theory, *instance-based learning theory* (IBLT; Gonzalez, Lerch, & Lebiere, 2003), can account for human behavior in a large diversity of tasks in which decisions are made from experience. Recently, they have demonstrated that the same computational model based on IBLT, without modifications, is able to account for multiple variations of the dual-choice paradigms commonly used to study decisions from experience (e.g., Gonzalez & Dutt, 2011; Lejarraga et al., 2012). In what follows, I present a general view of theories of learning in dynamic tasks. Then I expand on the IBLT as a general theory of decision making in dynamic tasks and on the cognitive models based on IBLT proposed to study decisions from experience in simple repeated-choice paradigms like those described earlier (e.g., Figure 13.3).

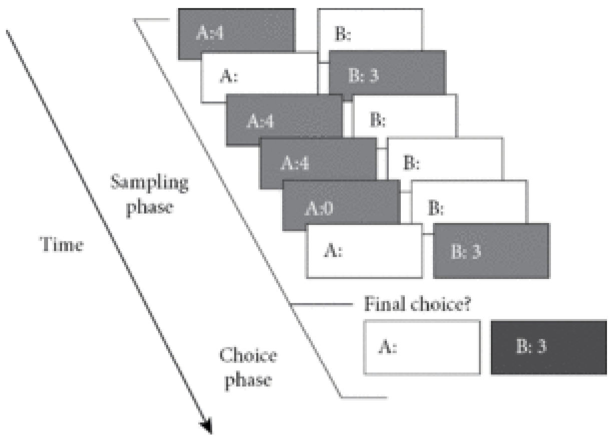

Click to view larger

Fig. 13.3 The sampling paradigm of decisions from experience.

Learning in DDM

Research on learning in and about dynamic systems indicates that humans remain suboptimal decision makers even after extended practice and after being given unlimited time and performance incentives (Diehl & Sterman, 1995; Sterman, 1994). That is, humans do not always improve their decision making from experience (Brehmer, 1980). One main impediment to learning in dynamic tasks is difficulty in processing feedback, particularly delayed feedback (Brehmer, 1992; Sterman, 1989*a*, 1989*b*). However, many other difficulties have also been documented, including our inabilities to deal with time constraints, high workload, and the limitations in our inherent cognitive abilities (Gonzalez, 2004, 2005*a*; Gonzalez, Thomas, & Vanyukov, 2005).

Various accounts have been proposed regarding how a human learns in dynamic systems (for summaries of these, see Busemeyer, 2002; Gonzalez, 2005*b*). One theory is that specific instances are used to control dynamic systems (Dienes & Fahey, 1995). This learning model was based on two cognitive mechanisms that compete every time someone encounters a decision-making situation: an algorithm and a set of context–action exemplars. The algorithm is a general heuristic or rule used in a novel situation, whereas the context–action exemplars are discrete representations of knowledge called "instances," a name derived from Logan's instance theory of automatization (1988). In this model, an implicit assumption is that a decision maker stores actions and their outcomes together in memory and retrieves them on the basis of their similarity to subsequently encountered situations.

Another theory of learning is proposed by the connectionist approach, in which decision making is built from interconnected units (Gibson, Fichman, & Plaut, 1997). This model is based on the control theory approach proposed by Brehmer (p. 255) (1990) and was implemented computationally via neural networks. It assumes that decision makers use outcome feedback to form two submodels: the *judgment submodel* that represents how the decision maker's actions affect outcomes and the *choice submodel* that represents which actions are taken to achieve desired outcomes. The judgment submodel learns by minimizing the differences between the outcomes it predicts and the outcomes received from feedback, whereas the *choice submodel* learns by minimizing the differences between the alternatives predicted by the judgment model and the alternatives actually selected. This model provides a good account of individuals' learning in dynamic situations and knowledge transfer to novel situations (Gibson, 2000).

IBLT, a third theory (Gonzalez et al., 2003; Gonzalez & Lebiere, 2005), was developed to reproduce decision-making behavior in dynamic tasks. It characterizes learning by storing in memory a sequence of action–outcome links produced by experienced events through a feedback loop process of human and environment interactions. Because of this theory's current relevance, not only for DDM but also for explaining decisions from experience (Gonzalez & Dutt, 2011), we expand on this theory and the cognitive model implementations of its mechanisms in the next section.

Generically speaking, the three learning models just summarized incorporate at least two common characteristics: 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 its structure (i.e., their knowledge is implicit), which suggests that the knowledge acquired was not in the form of rules about how the system 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 made decisions (Berry & Broadbent, 1987, 1988). Second, these models rely on a similarity process that determines the applicability of accumulated experiences to familiar situations. Research in analogical reasoning has demonstrated the increased relevance of analogy to learning and decision-making processes (Kurtz, Miao, & Gentner, 2001; Medin, Goldstone, & Markman, 1995). Decisions from experience are, very likely, the only method by which decisions are made in dynamic conditions. In fact, a recent study demonstrates that as a problem's complexity increases, people prefer to make decisions from experience rather than interpreting the given probabilities and outcomes of a one-shot decision (Lejarraga, 2010).

In summary, there are well-documented difficulties when humans make decisions in dynamic systems. Humans remain suboptimal or learn very slowly, often due to feedback delays, time constraints, and the cognitive workload required by these environments. To be able to understand and improve training protocols and guidelines, one needs to first understand how humans make decisions in these tasks. Fortunately, the similarities across the most prominent theories of learning in DDM converge on several issues. All models agree that humans learn facts—cause-and-effect knowledge related to the context—and none of the models presents the main form of learning as being structural knowledge or rules. Also, all of the models agree on the relevance of some form of recognition of familiar patterns from past experience; that is, that decisions are made from experience by retrieving a solution from similar past situations we've experienced.

IBLT

IBLT was developed to explain human decision-making behavior in dynamic tasks (Gonzalez et al., 2003). Dynamic decision making has been characterized by the process of making multiple, repeated, or sequential choices in conditions that evolve over time either as a result of previous decisions, with inaction, or spontaneously from the change occurring in the environment as time passes (Edwards, 1962).

Based on evidence from studies in naturalistic environments (Dreyfus & Dreyfus, 1986; Klein, Orasanu, Calderwood, & Zsambok, 1993; Pew & Mavor, 1998; Zsambok & Klein, 1997), laboratory studies with dynamic computer simulations (Microworlds) (Brehmer, 1990, 1992; Gonzalez, 2004, 2005*b*; Kerstholt & Raaijmakers, 1997), theoretical studies of decisions under uncertainty (Gilboa & Schmeidler, 1995, 2000), and other theories of learning in dynamic decision making (Dienes & Fahey, 1995; Gibson et al., 1997), IBLT proposed that decisions in dynamic tasks were made possible by referencing experiences from past similar situations and applying those decisions that worked in the past. IBLT's most important development was the description of the learning process and mechanisms by which experiences may be built, (p. 256) retrieved, evaluated, and reinforced during interactions with a dynamic environment.

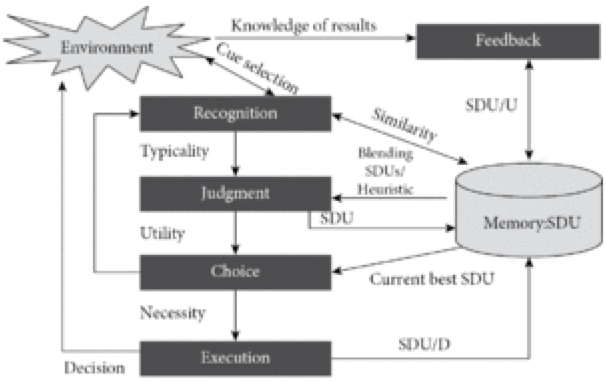

Click to view larger

Fig. 13.4 The IBLT process.

IBLT characterizes learning in dynamic tasks by storing "instances" in memory as a result of having experienced decision-making events. These instances are representations of three elements: a situation (S), which is defined by a set of attributes or cues; a decision (D), which corresponds to the action taken in situation S; and a utility or value (U), which is expected or received for making a decision D in situation S. IBLT proposes a generic decision-making process through which SDU instances are built, retrieved, evaluated, and reinforced (see a detailed description of this process in Gonzalez et al., 2003), with the steps consisting of recognition (similarity-based retrieval of past instances), judgment (evaluation of the expected utility of a decision in a situation through experience or heuristics), choice (decision on when to stop information search and select the optimal current alternative), execution (implementation of the decision selected), and feedback (update the utility of decision instances according to feedback) (see Figure 13.4).

When faced with a particular decision situation, people are likely to retrieve similar SDUs (SDUs with similar situations) from memory (Recognition step). In a typical situation (situation similar to past SDUs), the expected utility of an action is calculated by combining the utility of similar instances retrieved from memory (a procedure called *Blending*). In atypical situations, however, people fall back on heuristics in their evaluation of the expected utility of an action. Evaluation of a decision's expected utility in a given situation is done in the Judgment step. Alternative actions are evaluated sequentially, and, after each evaluation, the decision of whether more alternatives should be evaluated is determined by a *necessity mechanism*. Necessity may be subjectively determined by the decision maker's own preferences or by exogenous factors such as a lack of time or changes in the environmental conditions. The alternative with the highest utility among the evaluated alternatives is then selected (the Choice step) and executed (the Execution step), thus changing the environment and noting which SDU was executed in memory. Once a decision has been made, the decision's outcome is used as feedback to modify the utility value of the original SDUs (Feedback step).

The decision process of IBLT is determined by a set of learning mechanisms needed at different stages, including Blending (the aggregated weighted value of alternatives involving the instance's utility weighted by its probability of retrieval), Necessity (the decision to continue or stop exploring the environment), and Feedback (the selection of instances to be reinforced and the proportion by which the utility of these instances is reinforced).

In general, descriptive theories of behavior postulate processes and mechanisms that govern human behavior. IBLT proposed constructs and processes that rationalized the general phenomena of decision making in dynamic tasks and provided general explanations. These concepts and processes are generic and are motivated independently of a specific dynamic decision-making task or decision context. But it is the quantitative nature of a theory that can make it precise and testable. To test theories of human behavior, we use *computational* *(p. 257)* *models*: representations of some or all aspects of a theory as it applies to a particular task or context. Thus, the value of models is that they can solve concrete problems and provide explicit mathematical and computational representations of a theory, which can then be used to make predictions about behavior.

IBL Models

IBLT constructs and processes were implemented into a computational model that helped make the theory more explicit, transparent, and precise (Gonzalez et al., 2003). The first computational model based on IBLT (called Cog-IBLT) demonstrated the overall mechanisms and learning process in a dynamic and complex resource allocation task (reported in Gonzalez et al., 2003). Cog-IBLT was constructed within the ACT-R cognitive architecture (Anderson & Lebiere, 1998), using the cognitive mechanisms existing in ACT-R. Specifically, Cog-IBLT used the architecture's experimentally derived mathematical representations of *Activation* (a value that determines the usefulness of an instance from memory and experience and the relevance of the instance to the current context), *Partial Matching* (a value that determines the similarity of instances and the retrieval of instances that may be only similar to a current environmental situation), and *Retrieval Probability* (a value representing the probability of retrieving an instance as a function of Activation and Partial Matching). This model also used a modified version of the *Blending* concept proposed in Lebiere's dissertation (1998): an aggregate or combination of the values of multiple instances in memory. Through a series of "simulation experiments," Cog-IBLT demonstrated the explanatory and predictive potential of IBLT as it closely approximated the learning process from human data in a complex dynamic resource allocation task.

After the conceptualization of IBLT, many IBL models have been developed for a wide variety of tasks, including dynamically complex tasks (Gonzalez & Lebiere, 2005; Martin, Gonzalez, & Lebiere, 2004), training paradigms of simple and complex tasks (Gonzalez, Best, Healy, Kole, & Bourne, 2011; Gonzalez & Dutt, 2010), simple stimulus–response practice and skill acquisition tasks (Dutt, Yamaguchi, Gonzalez, & Proctor, 2009), and repeated binary-choice tasks (Lebiere, Gonzalez, & Martin, 2007; Lejarraga et al., 2012).

A recent IBL model has shown generalization across multiple tasks that share structural similarity with the paradigms used to study decisions from experience. Motivated by the work of Erev and Barron (2005), we built a model of repeated binary choice based on IBLT but within the ACT-R architecture (Lebiere et al., 2007). Erev and Barron (2005) demonstrated robust deviations from maximization in repeated binary choice and proposed the *reinforcement learning among cognitive strategies* (RELACS) model, which closely captures human data and outperforms other models. We argued for a simpler model, the IBL model, which was able to fit the data as well as RELACS (Lebiere et al., 2007).

The IBL model's development took an important turn when it was submitted to the Technion Prediction Tournament (TPT; Erev et al., 2010), a modeling competition that involved fitting and prediction phases in which the model authors were given a dataset to fit their models to and were evaluated in a novel dataset. The IBL model was developed independently and outside of ACT-R, and its mechanisms were isolated from all other ACT-R mechanisms (see Gonzalez, Dutt, & Lebiere, 2013, for a validation of this model within ACT-R and outside of ACT-R). Although this model did not win the TPT, its transparency, simplicity, and flexibility outside of ACT-R have been advantageous to recent theoretical developments. The model predicts performance better than the winner models of the TPT (Gonzalez & Dutt, 2011; Lejarraga et al., 2012) and predicts performance in a variety of repeated binary-choice tasks, probability-learning tasks, and dynamic-choice tasks across the multiple paradigms of decisions from experience and at the individual and team levels (Gonzalez & Dutt, 2011; Gonzalez, Dutt, & Lejarraga, 2011; Lejarraga et al., 2012). The discussions from this point on refer to this particular IBL model for repeated binary choice, which is explained in detail next.

The IBL Model of Repeated Binary Choice

Instances in a model of decision from experience paradigms (e.g., that shown in Figure 13.1) have a much simpler representation compared to instances in Cog-IBLT or in other IBL models. The instance structure is simple because the task structure is also simple. Each instance consists of a label that identifies a decision option in the task and the outcome obtained. For example, (Left, $4) is an instance in which the decision was to click the button on the left side, and the outcome obtained was $4. The details of this IBL model and its relevance were fully explained in (p. 258) Gonzalez and Dutt (2011), but its main aspects are summarized here.

The IBL model of repeated binary choice ("IBL model" hereafter) assumes that choices from experience are based on either repetition of past choices (i.e., "inertia") or on the aggregation of past experiences (i.e., "instances") of payoffs in memory that have been observed as a result of past choices (i.e., "blending"). At trial $t = 1$, the model starts with a random choice between the two options. Then, in each trial $t > 1$, the model first applies a probabilistic rule (based on a free parameter called *pInertia*) to determine whether or not to repeat its choice from the previous trial. If this probabilistic rule fails, then inertia does not determine the choice, and the model chooses the option with the highest *blended* value. An option's blended value is a weighted average of all observed payoffs on that option in previous trials. These observed payoffs are stored as instances in memory and are weighted such that the payoffs observed more frequently and recently receive a greater weight compared to infrequent and distant payoffs. This weight is a function of the recency and frequency of the instances' use, where the instance contains the observed payoffs. Formally, the model works as follows:

$$\text{In } t = 1 \quad \text{choose randomly between the two choice options} \qquad (1)$$

For each trial $t > 1$,

If the draw of a random value in the uniform distribution $U(0, 1) < pInertia$,

Then

Repeat the choice as made in the previous trial

Else

Select an option with the highest blended value as per Equation 2.

The blended value V of option j is:

$$V_j = \sum_{i=1}^{n} p_{ij} x_{ij} \quad (2)$$

where x_{ij} is the observed payoff in instance i for the option j, and p_{ij} is the probability of retrieving that instance for blending from memory (Gonzalez & Dutt, 2011; Lejarraga et al., 2012). Because the sampling paradigm involves a binary choice with two options, the values of j can be either 1 or 2 (i.e., right or left choice options). Thus, the blended value of an option j is the sum of all x_{ij} stored in instances in memory, weighted by their retrieval probability p_{ij}. The n value is the number of different instances containing the observed payoffs of option j up to the last trial. For example, if by trial t = 2, option j revealed 2 different payoffs stored in two instances, then n = 2 for option j. If the two observed payoffs on option j are the same in the previous two trials, then only one instance is created in memory and n = 1.

In any trial, the probability of retrieving an instance i containing a payoff observed for option j from memory is a function of that instance's activation relative to the activation of all other instances that contain observed payoffs l occurring within the same option. This probability is given by:

$$p_{ij} = \frac{e^{\frac{A_i}{\tau}}}{\sum_l e^{\frac{A_l}{\tau}}}, \quad (3)$$

where l refers to the total number of payoffs observed for option j up to the last trial, and τ is a noise value defined as $\sigma \cdot \sqrt{2}$ (Lebiere, 1998). The σ variable is a free noise parameter expected to capture the imprecision of recalling instances from memory from one trial to the next.

The activation of each instance in memory depends on the activation mechanism originally proposed in the ACT-R architecture (Anderson & Lebiere, 1998). The IBL model uses a simplified version of that activation mechanism. In each trial t, activation A of an instance i is

$$A_i = \ln\left[\sum_{t_i \in \{1,\ldots,t-1\}} (t - t_i)^{-d}\right] + \sigma \cdot \ln\left(\frac{1-\gamma_i}{\gamma_i}\right) \quad (4)$$

where d is a free decay parameter, and t_i refers to previous trials when the payoff contained in the instance i was observed (if a payoff occurs for the first time in a trial, a new instance containing this payoff is created in memory). The summation will include a number of terms that coincides with the number of times that a payoff has been observed after it was created (the time of creation of the instance itself is the first timestamp). Therefore, an instance's activation increases with the frequency of observing that payoff (i.e., by increasing the number of terms in the summation) and with the recency of observing that payoff (i.e., by small differences in (t – t_i)). The decay parameter d affects the activation of the instances directly because it captures the rate of forgetting. The higher the value of the *d* parameter, the faster the decay of instances' activations in memory.

The γ_i term is a random draw from a uniform distribution defined between 0 and 1, and (p. 259) $\sigma \cdot \ln\left(\frac{1-\gamma_i}{\gamma_i}\right)$ represents the Gaussian noise that is important for capturing variability in behavior from one trial to the next. The σ variable is the same noise parameter defined in Equation (3) above. A high σ implies greater noise in activation.

What the IBL Model Explains

The most recent developments of the IBL model are important given its simplicity and the broad predictions that it can make (e.g., Gonzalez & Dutt, 2011; Gonzalez et al., 2011; Lejarraga et al., 2012). Existing demonstrations from IBL models suggest the theory's generality and not only its descriptive power, but also its explanatory one. That is, the theory not only describes the kind of constructs and processes used in dynamic decision making, but it also helps explain why decision making in dynamic tasks occurs in the way described and not in other ways.

Two comprehensive and important demonstrations of the IBL model's robustness are the fitting and predicting of obtained results against a large and publicly available dataset, the TPT (Erev et al., 2010). TPT involved two types of experimental paradigms of decisions from experience—Sampling and Repeated choice—and all the problems in the TPT involved a choice between two options:

Safe: M with certainty

Risky: H with probability Ph; L otherwise (with probability 1−Ph)

A safe option offered a medium (M) payoff with certainty, and a risky option that offered a high (H) payoff with some probability (pH) and a low (L) payoff with the complementary probability. M, H, pH, and L were generated randomly, and a selection algorithm assured that the problems in each set differed in domain (positive, negative, and mixed payoffs) and probability (high, medium, and low pH).

The IBL model is able to predict the learning curves for most of the problems in the test set (see detailed tests in Lejarraga et al., 2012). The problems represent a large diversity of behavioral effects, and, in creating this diversity of problems, the TPT organizers (Erev et al., 2010) aimed at extending the traditional view of using counterexamples of particular behavioral effects by demonstrating the robustness of general learning effects. This demonstration and the additional ones found in Lejarraga et al. (2010) and in

Gonzalez and Dutt (2011) indicate the IBL model's ability to capture these general learning effects, too. Gonzalez (2012) discusses other effects that this model is able to capture and some phenomena that pose a challenge for the model. For example, IBL models explain payoff variability effect, underweighting of rare events, loss rate effect, individual differences, probability matching, and adaptation to nonstationary environments. But, in their current form, they are unable to capture the pure risk aversion effect; more risk seeking in losses compared to in gains domains; and emotions, social, and noncognitive effects. As we discuss in the last section of this chapter, future research is expected to address these and many other challenges that IBL models face.

Naturalistic Decision Making

In contrast to the traditional open-loop view and to the closed-loop view of decision making presented earlier, *naturalistic decision making* (NDM) is a field concerned with studying decision making in the "wild," in contexts where proficient decision makers draw conclusions from realistic cases and scenarios that are relevant to their experience and knowledge (Lipshitz, Klein, Orasanu, & Salas, 2001). Those who study naturalistic decision making are confronted with serious challenges. They often study large groups and real-life decision makers in complex decision situations. Real-world decision makers typically confront many uncertainties about the available options: they have inadequate information about their options, and they rarely know the likely costs and benefits or the value tradeoffs entailed. Although one could expect decision makers in the real world to have clear goals and to promote those goals with their decisions, the reality is that they are seldom rational, and, in fact, it is often hard to understand what constitutes rational choice under such conditions. Thus, these studies also have multiple limitations. Realistic studies demonstrate only particular examples of decision-making situations from which general predictions and inferences are hard to derive.

NDM is distinguishable from other views of decision making in two aspects: the research methods used (e.g., field observations rather than laboratory experimentation) and their focus on the study of experts in a particular context (e.g., firefighter professionals rather than university students). Methodologies used to gather information from experts are often based on observations and interviews, used in cognitive task analysis and in the detection of critical incidents or decisions (Lipshitz (p. 260) et al., 2001). These interviews and observations are conducted with people regarded as "experts" in a particular context. The most representative example of the NDM approach and findings is Gary Klein's work (1998). Through interviews and observations, Klein presents a study of firefighters. As a result of these studies, Klein proposed a model of decision making named *recognition-primed decision making* (RPD), which describes some of the processes involved in how experts make decisions in highly stressful and realistic conditions. This model includes the recognition of a situation as prototypical by the identification of relevant cues, the development of expectancies and suitable goals, and the identification of a typical course of action that can then be “played out” toward the future to be able to implement a course of action or generate further alternatives.

Although the NDM and the heuristics and biases traditions in decision-making research sharply contrast each other in their approaches and methods, leading researchers in both claim to have “failed to disagree” regarding the need to consider intuitive skill in every cognitive research effort (Kahneman & Klein, 2009; Kahneman, 2011). However, they also recognize that Daniel Kahneman “is still fascinated by persistent errors,” whereas Gary Klein “still recoils when biases are mentioned” (Kahneman & Klein, 2009, p. 525). We believe the DDM and closed-loop view of decision making presented in this chapter is one option pointing toward true reconciliation between these two extremes in decision-making research.

Conclusion

The field of behavioral decision research is broad and expanding, and it has a long and interesting history. This chapter attempted to provide a broad view of different topics and traditions in decision-making research where cognitive science in particular has been of the essence. The topics reviewed may be incomplete, but readers are encouraged to investigate the sources of the different themes to obtain a deeper understanding. The cognitive science perspective presented in this chapter is focused on the individual decision maker, and it attempts to highlight the importance of a closed-loop view in which cognitive processes and learning are critical. Given my interests and background in dynamic decision making and cognitive models of learning, this chapter provides slightly more weight to these themes. However, I expect that, in the future, more attention will be paid to the individual’s decision-making processes and to learning from the individual to the group and society levels. I also expect increased emphasis in complex, dynamic, and realistic environments. These contexts challenge many of the assumptions of current cognitive theories. A summary of these future directions is presented next.

Future Directions

I expect that, in the coming years, many questions and answers will emerge with respect to at least two aspects of decision making: decisions in dynamic environments and the connections from individual to small- and large-group decision making.

A lot of findings are accumulating with respect to simple, binary choice and monetary tasks, but I expect the current findings will expand to address more complex and dynamic decision situations. As the complexity of the tasks increases, so, too, will the need for additional cognitive mechanisms, such as those proposed by the ACT-R theory (Anderson & Lebiere, 1998) and those described in IBLT (Gonzalez et al., 2003). Complexity in dynamic tasks is defined not only by an increased number of options and more attributes, but also and most importantly by the interactions of these elements over time, a situation called *dynamic complexity*. Researchers have found that decision makers remain suboptimal even in the simplest dynamic system (with few options and attributes) after repeated practice, unlimited time, and performance incentives (Diehl & Sterman, 1995; Paich & Sterman, 1993; Sterman, 1989*a*, 1989*b*). Common causes include the multiple feedback processes, time delays, and nonlinearities involved in these systems and the "inability" to deal with such complexity (Cronin & Gonzalez, 2007; Cronin, Gonzalez, & Sterman, 2009). However, we need to understand and explain the underlying cognitive mechanisms leading to the learning difficulties in dynamic tasks (Gonzalez et al., 2003). The development of complex simulations for training purposes (e.g., for training the coordination of shipboard firefighting) opens up the possibility of bridging the gap separating laboratory studies of dynamic decision making from naturalistic studies of decision making.

Second, network science and complexity economics studies focus on the interactions among actors, decision makers, and their emergent social and economic phenomena, but they often oversimplify the cognitive aspects of the individuals involved. For example, to explain the complex dynamics seen in large economic systems like (p. 261) financial markets, researchers have often relied on agent-based models but rarely on *cognitive* models. On the other hand, cognitive modelers often focus on explaining individual behavior, relying on detailed cognitive models/architectures that formalize invariant cognitive representations and mechanisms (ACT-R, Anderson & Lebiere, 1998; Laird, Newell, & Rosenbloom, 1987), but they rarely model the behavior of a group of individuals (see Reitter & Lebiere, 2012, and Gonzalez, Dutt, & Lejarraga, 2011, for some exceptions). I expect that future research will bridge this gap by investigating the impact of individual cognitive characteristics on large groups and societies.

Many models of individual decisions from experience are incapable of representing human behavior in social contexts. For example, Erev and Roth (2001) noted that simple reinforcement learning models predicted the effect of experience in two-person games like the Iterated Prisoner's Dilemma (IPD) only in situations where players could not punish or reciprocate. A simple model predicts a decrease in cooperation over time, even though most behavioral experiments demonstrate an increase in mutual cooperation due to the possibility of reciprocation (Rapoport & Chammah, 1965; Rapoport & Mowshowitz, 1966). The IBL model appears to account for these reciprocity effects without the need for explicit and situation-specific rules (Gonzalez, Ben-Asher, Martin, & Dutt, 2015; Gonzalez et al., 2011). However, much work is needed to understand how the IBL model can be extended to account for effects at the levels of large groups and societies.

Acknowledgments

This chapter provides a broad overview of the cognitive perspective on decision making that I have built through years of research thanks to the support of a number of organizations including the Army Research Laboratory, the Office of Naval Research, the National Science Foundation, Defense Threat Reduction Agency, and the Army Research Office, among others; and thanks to my interactions with a number of doctoral students, postdoctoral fellows, and other research staff in the dynamic decision-making laboratory.

References

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum.

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). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology*, *79*(2), 251–272.

Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica*, *45*(1–3), 223–241.

Brehmer, B. (1990). Strategies in real-time, dynamic decision making. In R. M. Hogarth (Ed.), *Insights in decision making* (pp. 262–279). Chicago: University of Chicago Press.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, *81*(3), 211–241.

Busemeyer, J. R. (2002). Dynamic decision making. In N. J. Smelser & P. B. Baltes (Eds.), *International encyclopedia of the social and behavioral sciences* (Vol. 6, pp. 3903–3908). Oxford, UK: Elsevier Press.

Cronin, M., & Gonzalez, C. (2007). Understanding the building blocks of system dynamics. *System Dynamics Review*, *23*(1), 1–17.

Cronin, M., Gonzalez, C., & Sterman, J. D. (2009). Why don't well-educated adults understand accumulation? A challenge to researchers, educators and citizens. *Organizational Behavior and Human Decision Processes*, *108*(1), 116–130.

Crozier, R., & Ranyard, R. (1997). Cognitive process models and explanations of decision making. In R. Ranyard, W. R. Crozier, & O. Svenson (Eds.), *Decision making: Cognitive models and explanations* (pp. 5–20). New York: Routledge.

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.

Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine: The power of human intuition and expertise in the era of the computer*. New York: The Free Press.

Dutt, V., Yamaguchi, M., Gonzalez, C., & Proctor, R. W. (2009). An instance-based learning model of stimulus-response compatibility effects in mixed location-relevant and location-irrelevant tasks. In A. Howes, D. Peebles, & R. Cooper (Eds.), *Proceedings of the 9th International Conference on Cognitive Modeling—ICCM2009*. Manchester, UK.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors*, *4*(2), 59–73.

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology*, *32*, 53–88.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*(4), 912–931.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., et al. (2010). A choice prediction competition for choices from experience and from description. *Journal of Behavioral Decision Making*, *23*(1), 15–47.

Erev, I., & Roth, A. E. (2001). Simple reinforcement learning models and reciprocation in the Prisoner's Dilemma game. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 215–231). Cambridge, MA: MIT Press.

Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning* (pp. 182–189). New York: Academic Press.

Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review*, *83*(1), 37–64.

Flach, J. M., Hancock, P. A., Caird, J., & Vicente, K. J. (Eds.). (1995). *Global perspectives on the ecology of human-machine systems* (Vol. 1). Mahwah, NJ: Lawrence Erlbaum.

(p. 262) 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), 1–35.

Gigerenzer, G., Hertwig, R., & Pachur, T. (2011). *Heuristics: The foundation of adaptive behavior*. New York: Oxford University Press.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Gilboa, I., & Schmeidler, D. (1995). Case-based decision theory. *Quarterly Journal of Economics*, *110*(3), 605–639.

Gilboa, I., & Schmeidler, D. (2000). Case-based knowledge and induction. *IEEE Transactions on Systems, Man, and Cybernetics-PART A: Systems and Humans*, *30*(2), 85–95.

Goldstein, W. M., & Hogarth, R. M. (1997). Judgment and decision research: Some historical context. In W. M. Goldstein & R. M. Hogarth (Eds.), *Research on judgment and decision making* (pp. 3–65). New York: Cambridge University Press.

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. (2005*a*). The relationship between task workload and cognitive abilities in dynamic decision making. *Human Factors*, *47*(1), 92–101.

Gonzalez, C. (2005*b*). Decision support for real-time dynamic decision making tasks. *Organizational Behavior & Human Decision Processes*, *96*(2), 142–154.

Gonzalez, C. (2012). Training decisions from experience with decision making games. In P. Durlach & A. M. Lesgold (Eds.), *Adaptive technologies for training and education* (pp. 167–178). New York: Cambridge University Press.

Gonzalez, C., Ben-Asher, N., Martin, J. & Dutt, V. (2015). A Cognitive Model of Dynamic Cooperation with Varied Interdependency Information. *Cognitive Science, 39*, 457–495.

Gonzalez, C., Best, B. J., Healy, A. F., Kole, J. A., & Bourne, L. E., Jr. (2011). A cognitive modeling account of simultaneous learning and fatigue effects. *Journal of Cognitive Systems Research*, *12*(1), 19–32.

Gonzalez, C., & Dutt, V. (2010). Instance-based learning models of training. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *54*(27), 2319–2323.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating decisions from experience in sampling and repeated choice paradigms. *Psychological Review*, *118*(4), 523–551.

Gonzalez, C., Dutt, V., & Lebiere, C. (2013). Validating Instance-Based Learning Mechanisms Outside of ACT-R. *Journal of Computational Science*, 4, pp. 262–268. DOI information: 10.1016/j.jocs.2011.12.001.

Gonzalez, C., Dutt, V., & Lejarraga, T. (2011). A loser can be a winner: Comparison of two instance-based learning models in a market entry competition. *Games*, *2*(1), 136–162.

Gonzalez, C., & Lebiere, C. (2005). Instance-based cognitive models of decision making. In D. Zizzo & A. Courakis (Eds.), *Transfer of knowledge in economic decision-making* (pp. 148–165). New York: Macmillan (Palgrave Macmillan).

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591–635.

Gonzalez, C., Thomas, R. P., & Vanyukov, P. (2005). The relationships between cognitive ability and dynamic decision making. *Intelligence*, *33*(2), 169–186.

Griffin, D. W., Gonzalez, R., Koehler, D. J., & Gilovich, T. (2012). Judgmental heuristics: A historical overview. In K. J. Holyoak & R. G. Morrison (Eds.), *The Oxford handbook of thinking and reasoning* (pp. 322–345). New York: Oxford University Press.

Hastie, R. (2001). Problems for judgment and decision making. *Annual Review of Psychology*, *52*, 653–683.

Hertwig, R. (2012). Tapping into the wisdom of the crowd—with confidence. *Science*, *336*, 303–304.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, *64*(6), 515–526.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge: Cambridge University Press.

Kerstholt, J. H., & Raaijmakers, J. G. W. (1997). Decision making in dynamic task environments. In R. Ranyard, W. R. Crozier, & O. Svenson (Eds.), *Decision making: Cognitive models and explanations* (pp. 205–217). London: Routledge.

Klein, G. A. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.

Klein, G. A., Orasanu, J., Calderwood, R., & Zsambok, C. E. (Eds.). (1993). *Decision making in action: Models and methods*. Norwood, NJ: Ablex.

Kurtz, K. J., Miao, C., & Gentner, D. (2001). Learning by analogical bootstrapping. *Journal of Learning Sciences*, *10*(4), 417–446.

Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, *33*(1), 1–64.

Lebiere, C. (1998). *The dynamics of cognition: An ACT-R model of cognitive arithmetic* (Doctoral dissertation). Retrieved from **http://reports-archive.adm.cs.cmu.edu/**. (Technical Report CMU-CS-98-186)

Lebiere, C., Gonzalez, C., & Martin, M. (2007). Instance-based decision making model of repeated binary choice. In R. L. Lewis, T. A. Polk, & J. E. Laird (Eds.), *Proceedings of the 8th International Conference on Cognitive Modeling* (pp. 67–72). Oxford, UK: Psychology Press.

Lee, W. (1971). *Decision theory and human behavior*. Oxford, UK: John Wiley & Sons.

Lejarraga, T. (2010). When experience is better than description: Time delays and complexity. *Journal of Behavioral Decision Making*, *23*(1), 100–116.

Lejarraga, T., Dutt, V., & Gonzalez, C. (2012). Instance-based learning: A general model of repeated binary choice. *Journal of Behavioral Decision Making*, *25*(2), 143–153.

(p. 263) Lipshitz, R., Klein, G., Orasanu, J., & Salas, E. (2001). Focus article: Taking stock of naturalistic decision making. *Journal of Behavioral Decision Making*, *14*(5), 331–352.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*(4), 492–527.

Martin, M. K., Gonzalez, C., & Lebiere, C. (2004). Learning to make decisions in dynamic environments: ACT-R plays the beer game. In M. C. Lovett, C. D. Schunn, C. Lebiere, & P. Munro (Eds.), *Proceedings of the Sixth International Conference on Cognitive Modeling* (Vol. 420, pp. 178–183). Pittsburgh, PA: Lawrence Erlbaum.

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.

Paich, M., & Sterman, J. D. (1993). Boom, bust and failures to learn in experimental markets. *Management Science*, *39*(12), 1439–1458.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.

Pew, R. W., & Baron, S. (1978). The components of an information processing theory of skilled performance based on optimal control perspective. In G. E. Stelmach (Ed.), *Information processing in motor control and learning* (pp. 71–78). New York: Academic Press.

Pew, R. W., & Mavor, A. S. (1998). *Modeling human and organizational behavior*. Washington, DC: National City Press.

Rapoport, A., & Chammah, A. M. (1965). *Prisoner's dilemma: A study in conflict and cooperation*. Ann Arbor: University of Michigan Press.

Rapoport, A., & Mowshowitz, A. (1966). Experimental studies of stochastic models for the Prisoner's dilemma. *System Research and Behavioral Science*, *11*(6), 444–458.

Reitter, D., & Lebiere, C. (2012). Social cognition: Memory decay and adaptive information filtering for robust information maintenance. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence* (pp. 242–248). Toronto, Canada: AAAI.

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, *15*(3), 233–250.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*(1), 99–118.

Simon, H. A. (1957). *Models of man: Social and rational*. New York: Wiley.

Sterman, J. D. (1989*a*). Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes*, *43*(3), 301–335.

Sterman, J. D. (1989*b*). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science*, *35*(3), 321–339.

Sterman, J. D. (1994). Learning in and about complex systems. *System Dynamics Review*, *10*(2–3), 291–330.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.

Wickens, C. D., & Kramer, A. (1985). Engineering psychology. *Annual Review of Psychology*, *36*, 307–348.

Zsambok, C. E., & Klein, G. A. (Eds.). (1997). *Naturalistic decision making*. Mahwah, NJ: Lawrence Erlbaum. **(p. 264)**

Cleotilde Gonzalez

Cleotilde Gonzalez, Dynamic Decision Making Laboratory, Social and Decision Sciences Department, Carnegie Mellon University