CHAPTER 84

Decision-Making Models

MARK R. LEHTO
Purdue University

Handbook of Industrial Engineering: Technology and Operations Management, Third Edition.
Edited by Gavriel Salvendy

1. INTRODUCTION

This chapter focuses on the broad topic of human decision making. Decision making is often viewed as a stage of human information processing because people must gather, organize, and combine information from different sources to make decisions. However, as decisions grow more complex, information processing actually becomes part of decision making and methods of decision support that help decision makers process information become of growing importance. Decision making also overlaps with problem solving. The point where decision making becomes problem solving is fuzzy, but many decisions require problem solving, and the opposite is true as well. Cognitive models of problem solving are consequently relevant for describing many aspects of human decision making. They become especially relevant for describing steps taken in the early stages of a decision where choices are formulated and alternatives are identified.

A complete treatment of human decision making is well beyond the scope of a single book chapter.* The topic has its roots in economics and is currently a focus of operations research and management science, psychology, sociology, and cognitive engineering. These fields have produced numerous models and a substantial body of research on human decision making. At least three objectives have motivated this work: to develop normative prescriptions that can guide decision makers, to describe how people make decisions and compare the results to normative prescriptions, and to determine how to help people apply their ‘‘natural’’ decision-making methods more successfully. The goals of this chapter are to synthesize the elements of this work into a single picture and to provide some depth of coverage in particularly important areas. The integrative model presented in Section 1.3 focuses on the first goal. The remaining sections address the second goal.

1.1. Role and Utility of Chapter

This chapter is intended to provide an overall perspective on human decision making to human factors practitioners, developers of decision tools (such as expert systems), product designers, researchers in

*No single book covers all of the topics addressed here. More detailed sources of information are references throughout the chapter. Sources such as von Neuman and Morgenstern (1947), Friedman (1990), Savage (1954), Luce and Raiffa (1957), and Shafer (1976), are useful texts for people desiring an introduction to normative decision theory. Raiffa (1968), Keeney and Raiffa (1976), Saaty (1988), Buck (1989), and Clemen (1996) provide applied texts on decision analysis. Kahneman et al. (1982), Winterfeldt and Edwards (1986), Sevenson and Maule (1993), Payne et al. (1993), Yates (1994), and Heath et al. (1994), among numerous others, provide recent texts addressing elements of behavioral decision theory. Klein et al. (1993) and Klein (1998) provide introductions to naturalistic decision making.

related areas, and others who are interested in both how people make decisions and how decision making might be improved. The chapter consequently presents a broad set of prescriptive and descriptive approaches. Numerous applications are presented and strengths and weaknesses of particular approaches are noted. Emphasis is also placed on providing useful references containing additional information on topics the reader may find of special interest.

Section 2 addresses topics grouped under the somewhat arbitrary heading of classical decision theory. The presented material provides a normative and prescriptive framework for making decisions. Section 3 summarizes decision analysis, or the application of normative decision theory to improve decisions. The discussion considers the advantages of the various approaches, how they can be applied, and what problems might arise during their application. Section 4 addresses topics grouped under the heading of behavioral decision theory. The material in the latter section compares human decision making to the normative models discussed earlier. Several descriptive models of human judgment, preference, and choice are also discussed. Section 5 explores topics falling under the heading of dynamic and naturalistic decision theory. This material should be of interest to practitioners interested in the process followed when many real-world decisions are made, the quality of these decisions, and why people use particular methods to make decisions. The discussion provides insight into how people perform diagnostic tasks, make decisions to take risks when using products, and develop expertise. Section 6 introduces the topic of group decision making. The discussion addresses conflict resolution both within and between groups, group performance and biases, and methods of group decision making.

1.2. Elements of Decision Making

Decision making requires that the decision maker make a choice between two or more alternatives (note that doing nothing can be viewed as making a choice). The selected alternative then results in some real or imaginary consequences to the decision maker. Judgment is a closely related process where a person rates or assigns values to attributes of the considered alternatives. For example, a person might judge both the safety and attractiveness of a car being considered for purchase. Obtaining an attractive car is a desirable consequence of the decision, while obtaining an unsafe car is an undesirable consequence. A rational decision maker seeks desirable consequences and attempts to avoid undesirable consequences.

The nature of decision making can vary greatly, depending on the decision context. Certain decisions, such as deciding where and what to eat for lunch, are routine and repeated often. Other choices, such as purchasing a house, choosing a spouse, or selecting a form of medical treatment for a serious disease, occur seldom, may involve much deliberation, and take place over a longer time period. Decisions may also be required under severe time pressure and involve potentially catastrophic consequences, such as when a fire chief decides whether to send fire fighters into a burning building. Previous choices may constrain or otherwise influence subsequent choices (for example, a decision to enter graduate school might constrain a future employment-related decision to particular job types and locations). The outcomes of choices may be uncertain and in certain instances are determined by the actions of potentially adverse parties, such as competing manufacturers of a similar product. Decisions may be made by a single individual or by a group. Within a group, there may be conflicting opinions and differing degrees of power between individuals or factions. Decision makers may also vary greatly in their knowledge and degree of aversion to risk.

Conflict occurs when a single decision maker is not sure which choice should be selected or when there is lack of consensus within a group regarding the choice. Both for groups and single decision makers, conflict occurs, at the most fundamental level, because of uncertainty or conflicting objectives. Uncertainty can take many forms and is one of the primary reasons decisions can be difficult. In ill-structured decisions, decision makers may not have identified the current condition, alternatives to choose between, or their consequences. Decision makers also may be unsure what their aspirations or objectives are, or how to choose between alternatives. After a decision has been structured, at least four reasons for conflict may exist. First, when alternatives have both undesirable and desirable consequences, decision makers may experience conflict due to conflicting objectives. For example, a decision maker considering the purchase of an air bag-equipped car may experience conflict because an air bag increases cost as well as safety. Second, decision makers may be unsure of their reaction to a consequence. For example, people considering whether to enter a raffle where the prize is a sailboat may be unsure how much they want a sailboat. Third, decision makers may not know whether a consequence will happen for sure. Even worse, they may be unsure what the probability of the consequences is, or may not have enough time to evaluate the situation carefully. They also may be uncertain about the reliability of information they have. For example, it may be difficult to determine the truth of a sale person's claim regarding the probability of their product breaking down immediately after the warranty expires.

To resolve conflicts, decision makers must deal appropriately with uncertainty, conflicting objectives, or a lack of consensus. Conflict resolution, therefore, becomes a primary focus of decision theory. The following section presents an integrative model of decision making that relates conflict resolution to the above-discussed elements of decision making. This model specifically considers how decision making changes when different sources of conflict are present. It also matches methods of conflict resolution to particular sources of conflict and decision rules.

1.3. Integrative Model of Decision Making

Human decision making can be viewed as a stage of information processing that falls between perception and response execution (Welford 1976). The integrative model of human decision making, presented in Figure 1, shows how the elements of decision making discussed above fit into this perspective. From this view, decision making is the process followed when a response to a perceived

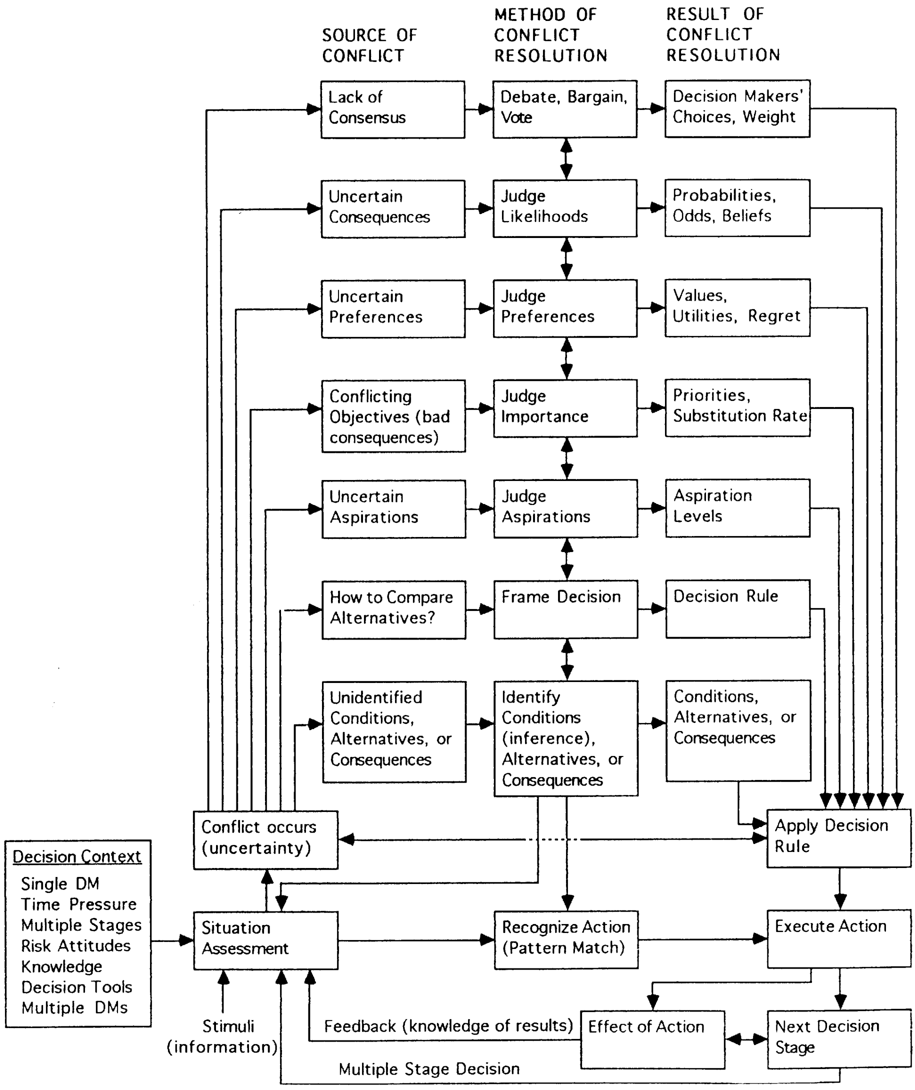

Figure 1 Integrative Model of Human Decision Making.

stimulus is chosen. The process followed depends on what decision strategy is applied and can vary greatly between decision contexts.* Decision strategies, in Figure 1, correspond to different paths between situation assessment and executing an action. The particular decision strategy followed depends upon both the decision context and whether or not the decision maker experiences conflict.†

At least four, sometimes overlapping, categories of decision making can be distinguished. *Group decision making* occurs when multiple decision makers interact and is represented at the highest level of the model as a source of conflict that might be resolved through debate, bargaining, or voting. For example, members of a university faculty committee might debate and bargain before voting between candidates for a job opening.

Dynamic decision making occurs in a changing environment, in which the results of earlier decisions impact future decisions. The decisions made in such settings often make use of feedback and are multistage in nature. For example, a decision to take a medical test almost always requires a subsequent decision regarding what to do after receiving the test results. Dynamic decision making is represented at the lowest level of the model by the presence of two feedback loops, which show how the action taken and its effects can feedforward to the assessment of a new decision or feedback to the reassessment of the current decision.

Routine decision making occurs when decision makers use knowledge and past experience to decide quickly what to do and is especially prevalent in dynamic decision making contexts. Routine decision making is represented in Figure 1 as a single pattern-matching step or associative leap between situation assessment and executing an action. For example, a driver after perceiving a stop sign decides to stop, or the user of a word-processing system after perceiving a misspelled word decides to activate the spell checker. Since routine decisions are often made in dynamic task environments, routine decision making is discussed in this chapter as a subtopic of dynamic decision making.

Conflict-driven decision making occurs when various forms of conflict must be resolved before an alternative action can chosen and often involves a complicated path between situation assessment and executing an action.‡ Before executing an action, the decision maker experiences conflict, somehow resolves it, and then either recognizes the best action (conflict resolution might transform the decision to a routine one) or applies a decision rule. Applying the decision rule ideally leads to a choice that is then executed. Attempting to apply the decision rule may, however, cause additional conflicts, leading to more conflict resolution. For example, decision makers may realize they need more information to apply a particular decision rule. In response, they might decide to use a different decision rule that requires less information. Along these lines, when choosing a home, a decision maker might decide to use a satisficing decision rule after seeing that hundreds of homes are listed in the classified ads of the local newspaper.

Potential sources of conflict, methods of conflict resolution, and the results of conflict resolution are listed at the top of Figure 1. Each source of conflict maps to a particular method of conflict resolution, which then provides a result necessary to apply a decision rule, as schematically illustrated in the figure.§ Table 1 presents a set of decision rules, briefly describes their procedural nature and their required inputs, and also lists the sections of this chapter where they are covered. The required inputs of particular decision rules can be easily mapped to sources of conflict. As shown in the table, each decision rule requires that alternatives and their consequences be identified. Other decision rules require measures of aspiration, importance, preference, and uncertainty for each consequence or consequence dimension. For example, to compare alternatives using expected value, the probability and value of each consequence must also be known. Certain decision rules also accept inputs describing the degree of consensus between decision makers.

Accordingly, conflict occurs at the most fundamental level when the current condition, alternative actions, or their consequences have not been identified. At the next most fundamental level, conflict occurs when the decision maker is unsure how to compare the alternatives. In other words, the

* The notion that the best decision strategy varies between decision contexts is a fundamental assumption of the theory of contigent decision making (Payne et al. 1993), cognitive continuum theory (Hammond 1980), and other approaches discussed later in this chapter.

† Conflict has been recognized as an important determinant of what people will do in risky decision-making contexts (Janis and Mann 1977). Janis and Mann focus on the stressful nature of conflict and on how affective reactions in stressful situations can impact the decision strategies followed.

‡ The distinction between routine and conflict-driven decision making made here is similar to Rasmussen's (1983) distinction between (a) routine skill or rule-based levels of control and (b) nonroutine knowledge-based levels of control in information-processing tasks.

§ Note that multiple sources of conflict are possible for a given decision context. An attempt to resolve one source of conflict may also make the decision maker aware of other conflicts that must first be resolved. For example, decision makers may realize they need to know what the alternatives are before they can determine their aspiration levels.

TABLE 1 Decision Rules, Required Inputs, and Procedure Applied by the Rules

Decision Rule	Required Inputs	Procedure Applied	Section Covered
Dominance	All alternatives, value of each consequence.	Select alternative best on all consequences.	2.1.2
EBA	All alternatives, value of each consequence.	Select first alternative found to be best on a consequence dimension. Random order of consequences.	2.1.3
Lexicographic	All alternatives, value of each consequence, priorities.	Order consequences by priority. Select first alternative found to be best on a consequence dimension.	2.1.3
Satisficing	At least one and up to all alternatives, aspiration level and value of each consequence	Sequentially evaluate each alternative. Stop if each consequence of an alternative equals or exceeds the aspiration level.	2.1.4
Minimax Cost	All alternatives, value of each consequence.	Compare the worst consequence values of each alternative	2.1.5
Minimax Regret	All alternatives, regret for each consequence.	Compare largest regrets of each alternative	2.1.5
EV	All alternatives, probability and value of each consequence.	Weight value of each consequence by its probability, for each alternative.	2.1.6
Laplace	All alternatives, value or utility of each consequence.	Weight value or utility of each consequence equally, for each alternative.	2.1.7
SEU	All alternatives, probability and utility of each consequence.	Weight utility of each consequence by its probability, for each alternative.	2.1.7
MAUT	All alternatives, value or utility of each consequence, priorities.	Weight value or utility of each consequence by priority, for each alternative.	2.1.8
Holistic	All alternatives and consequences.	Holistically compare the consequences of each alternative	2.1.9

decision maker has not yet selected a decision rule. Given that the decision maker has a decision rule, conflict can still occur if the needed inputs are not available. These sources of conflict and associated methods of conflict resolution are briefly addressed below in relation to the remainder of this chapter.

Identifying the current condition, alternative actions, and their consequences is an important part of decision making. This topic is emphasized in both naturalistic decision theory (Klein et al. 1993) and decision analysis* (Raiffa 1968; Clemen 1996). Decision trees, influence diagrams, and other tools for structuring decisions are covered in Section 3. Normative methods of identifying the current condition falling under the topic of inference (or diagnosis) are presented in Section 2.2. Section 4.1 describes several descriptive models of human inference and discusses their limitations. Section 6 includes discussion group decision-making methods that may be useful at this decision-making stage.

When decision makers are unsure how to compare alternatives, they must consider what information is available and then frame the decision appropriately. The way the decision is framed then determines (1) which decision rules are appropriate, (2) what information is needed to make the decision using the given rules (as discussed earlier in reference to Table 1), and (3) the choices selected. As discussed in Section 5.1, there is reason to believe that people apply different decision-making strategies in different decision contexts. Section 2.1 discusses the appropriateness of decision

*Clemen (1996) includes a chapter on creativity and decision structuring. Some practitioners claim that structuring the decision is the greatest contribution of the decision analysis process.

rules and how the particular rule used can impact choices. When the specific inputs needed by a decision rule are not available, the resulting conflict might be resolved by judging aspirations, importance, preference, or likelihood. It also might be resolved by choosing a different decision rule or strategy. As noted in Section 5.1, there is a prevalent tendency among decision makers in naturalistic settings to minimize analysis and its required cognitive effort. In group situations, conflict due to a lack of consensus between multiple decision makers might be resolved through debate, bargaining, or voting (Section 6).

2. CLASSICAL DECISION THEORY

Classical decision theory began with the development of normative models in economics and statistics that specified optimal decisions (von Neumann and Morgenstern 1947; Savage 1954). Classical decision theory focuses heavily on the notion of rationality (Winterfeldt and Edwards 1986; Savage 1954). Emphasis is placed on the quality of the process followed when making a decision rather than on the ultimate outcome. Accordingly, a rational decision maker must think logically about the decision. To do this, the decision maker must first formally describe what is known about the decision. The decision is then made by applying principles of logic and Bayesian probability theory (Savage 1954). This approach is therefore quantitative, and also normative or prescriptive if the numerical inputs needed are available.

The classical approach has been applied to two related problems: (1) preference and choice, and (2) statistical inference.

2.1. Choice Procedures

Classical decision theory represents preference and choice problems in terms of four basic elements: (1) a set of potential actions (A_i) to choose between, (2) a set of events or world states (E_j), (3) a set of consequences (C_{ij}) obtained for each combination of action and event, and (4) a set of probabilities (P_{ij}) for each combination of action and event. For example, a decision maker might be deciding whether to wear a seatbelt when traveling in an automobile. Wearing or not wearing a seat belt corresponds to two actions A_1 and A_2. The expected consequence of either action depends upon whether an accident occurs. Having or not having an accident corresponds to two events E_1 and E_2. Wearing a seatbelt reduces the expected consequences C_{11} having an accident (E_1). As the probability of having an accident increases, use of a belt should therefore become more attractive.

Once a decision has been represented in terms of these basic elements, the choice is then made by applying decision rules. Numerous decision rules have been developed. Decision rules are based upon basic axioms (or what are felt to be self-evident assumptions) of rational choice. Not all rules, however, make use of the same axioms. Different rules make different assumptions and can provide different preference orderings for the same basic decision. The following discussion will first present some of the most basic axioms. Then several well-known decision rules will be briefly covered.

2.1.1. Axioms of Rational Choice

Numerous axioms have been proposed that are essential either for a particular model of choice or for the method of eliciting numbers used for a particular model (Winterfeldt and Edwards 1986). The best-known set of axioms (Table 2) establishes the normative principle of subjective expected utility (SEU) as a basis for making decisions (see Savage 1954; Luce and Raiffa 1957 for a more rigorous description of the axioms). On an individual basis, these axioms are intuitively appealing (Stukey and Zeckhauser 1978), but, as discussed in Section 4, people's preferences can deviate significantly from the SEU model in ways that conflict with certain axioms. Consequently, there has been a movement toward developing other, less restrictive standards of normative decision making (Frisch and Clemen 1994; Zey 1992).

Frisch and Clemen propose that "a good decision should (a) be based on the relevant consequences of the different options (*consequentialism*), (b) be based on an accurate assessment of the world and a consideration of all relevant consequences (*thorough structuring)*, and (c) make tradeoffs of some form (*compensatory decision rule*)." Consequentialism and the need for thorough structuring are both assumed by all normative decision rules. Most normative rules are also compensatory. However, when people make routine habitual decisions, they often don't consider the consequences of their choices, as discussed in Section 5. Also, because of cognitive limitations and the difficulty of obtaining information, it becomes unrealistic in many settings for the decision maker to consider all the options and possible consequences. To make a decision under such conditions, decision makers may limit the scope of the analysis by applying principles such as satisficing and other noncompensatory decision rules discussed below. They also may apply heuristics, based on their knowledge or experience, leading to performance that can approximate the results of applying compensatory decision rules (Section 4).

TABLE 2 Basic Axioms of Subjective Expected Utility Theory

A. Ordering/Quantification of Preference
Preferences of decision makers between alternatives can be quantified and ordered using the relations: $>$, where $A > B$ means that A is preferred to B $=$, where $A = B$ means that A and B are equivalent $\geq$, where $A \geq B$ means that B is not preferred to A
B. Transitivity of Preference
if $A_1 \geq A_2$ and $A_2 \geq A_3$, then $A_1 \geq A_3$
C. Quantification of Judgment
The relative likelihood of each possible consequence that might result from an alternative action can be specified.
D. Comparison of Alternatives
If two alternatives yield the same consequences, the alternative yielding the greater chance of the preferred consequence is preferred.
E. Substitution
If $A_1 > A_2 > A_3$, then the decision maker will be willing to accept a gamble $[p(A_1)$ and $(1 - p)(A_3)]$ as a substitute for A_2 for some value of $p \geq 0$.
F. Sure Thing Principle
If $A_1 \geq A_2$, then for all p, the gamble $[p(A_1)$ and $(1 - p)(A_3)] \geq [p(A_2)$ and $(1 - p)(A_3)]$.

2.1.2. Dominance

Dominance is perhaps the most fundamental normative decision rule. Dominance is said to occur between two alternative actions A_i and A_j when A_i is at least as good as A_j for all events E, and for at least one event E_k, A_i is preferred to A_j. For example, one investment might yield a better return than another regardless of whether the stock market goes up or down. Dominance can also be described for the case where the consequences are multidimensional. This occurs when for all events E, the kth consequence associated with action i (C_{ik}) and action j (C_{jk}), satisfies the relation $C_{ik} \geq C_{jk}$ for all k, and for at least one consequence $C_{ik} > C_{jk}$. For example, a physician choosing between alternative treatments has an easy decision if one treatment is *both* cheaper and more effective for all patients.

Dominance is obviously a normative decision rule, since a dominated alternative can never be better than the alternative that dominates it. Dominance is also conceptually simple, but it can be difficult to detect when there are many alternatives to consider or many possible consequences. The use of tests for dominance by decision makers in naturalistic settings in discussed further in Section 5.1.5.

2.1.3. Lexicographic Ordering and EBA

The lexicographic ordering principle (see Fishburn 1974) considers the case where alternatives have multiple consequences. For example, a purchasing decision might be based on both the cost and performance of the considered product. The different consequences are first ordered in terms of their importance. Returning to the above example, performance might be considered more important than cost. The decision maker then sequentially compares each alternative beginning with the most important consequence. If an alternative is found that is better than the others on the first consequence, it is immediately selected. If no alternative is best on the first dimension, the alternatives are compared for the next-most important consequence. This process continues until an alternative is selected or all the consequences have been considered without making a choice. The latter situation can happen only if the alternatives have the same consequences.

The elimination by aspects (EBA) rule (Tversky 1972) is similar to the lexicographic decision rule. It differs in that the consequences used to compare the alternatives are selected in random order, where the probability of selecting a consequence dimension is proportional to its importance. Both EBA and lexicographic ordering are noncompensatory decision rules, since the decision is made using a single consequence dimension. Returning to the above example, the lexicographic principle would result in selecting a product with slightly better performance, even if it costs much more. EBA would select either product depending on which of the consequences was first selected.

2.1.4. Minimum Aspiration Level and Satisficing

The minimum aspiration level or satisficing decision rule assumes that the decision maker sequentially screens the alternative actions until an action is found which is good enough. For example, a person considering the purchase of a car might stop looking once he or she found an attractive deal instead of comparing every model on the market. More formally, the comparison of alternatives stops once a choice is found that exceeds a minimum aspiration level S_{ik} for each of its consequences C_{ik} over the possible events E_k.

Satisficing can be a normative decision rule when (1) the expected benefit of exceeding the aspiration level is small, (2) the cost of evaluating alternatives is high, or (3) the cost of finding new alternatives is high. More often, however, it is viewed as an alternative to maximizing decision rules. From this view, people cope with incomplete or uncertain information and their limited rationality by satisficing in many settings instead of optimizing (Simon 1955, 1983).

2.1.5. Minimax (Cost and Regret) and the Value of Information

Minimax cost selects the best alternative (A_i) by first identifying the worst possible outcome for each alternative. The worst outcomes are then compared between alternatives. The alternative with the minimum worst-case cost is selected. Formally, the preferred action A_i is the action for which over the events k, $\text{MAX}_k(C_{ik}) = \text{MIN}_i[\text{MAX}_k(C_{ik})]$. For example, in Table 3, the maximum cost is 5 for alternative A_1, 7 for A_2, and 8 for A_3. A_1 would be chosen since it has the smallest maximum cost. Minimax cost corresponds to assuming the worst and therefore makes sense as a strategy where an adverse opponent is able to control the events (von Neumann and Morgenstern 1947). Along these lines, an airline executive considering whether to reduce fares might assume that a competitor will also cut prices, leading to a no-win situation.

Minimax regret involves a similar process, but the calculations are performed using regret instead of cost (Savage 1954). Regret is calculated by first identifying which alternative is best for each possible event. The regret R_{ik}, associated with each consequence (C_{ik}) for the combination of event E_k and alternative A_i then becomes: $R_{ik} = \text{MAX}_i(C_{ik}) - C_{ik}$. Returning to our earlier example, if E_1 occurs, alternative A_2 with a cost of 2 is best, resulting in a regret of 0 ($2 - 2$). A_1 has a cost of 5, resulting in a regret of 3 ($5 - 2$). A_3 has a cost of 6, resulting in a regret of 4 ($6 - 2$). These calculations are repeated for events E_2 and E_3, resulting in regret values for each combination of events and alternative actions. The preferred action A_i is the action for which over the events k, $\text{MAX}_k(R_{ik}) = \text{MIN}_i[\text{MAX}_k(R_{ik})]$. The maximum regrets for A_1 (a value of 3) and A_3 (a value of 4) are both found when event E_1 occurs. The maximum regret for A_2 (a value of 2) is found when event E_2 occurs. Alternative A_2 is then selected because it has the minimum maximum regret.

Note that the minimax cost and minimax regret principles do not always suggest the same choice (Table 3). Minimax cost is easily interpreted as a conservative strategy. Minimax regret is more difficult to judge from an objective or normative perspective (Savage 1954). As shown by the example, minimax regret can be less conservative than minimax cost. Alternatives that were not chosen can also impact choices made using minimax regret. For example, if alternative A_3 is removed from consideration, minimax regret and minimax cost will both select A_1. The interesting conclusion is that comparative and absolute measures of preference can result in different choices.

Bell (1982) argues persuasively that regret plays a very prominent role in decision making under uncertainty. For example, the purchaser of a new car might be happy, until finding out that a neighbor got the same car for $200 less from a different dealer. It is interesting to observe that regret is closely related to the value of information. This follows, since with hindsight, decision makers may regret their choice if they did not select the alternative giving the best result for the event (E_k) which actually took place. With perfect information, the decision maker would have chosen E_k. Conse-

TABLE 3 Example Comparison of Minimax Cost and Minimax Regret. Minimax cost selects A_1 and minimax regret selects A_2

	E_1	E_2	E_3	Max Cost	Max Regret
A_1	5	5	5	**5**	3
A_2	2	7	2	7	**2**
A_3	6	8	4	8	4

quently, the regret (R_{ik}) associated with having chosen alternative (A_i) is a measure of the value of having perfect information, or of knowing ahead of time that event E_k would occur. When each of the events (E_k) occur with probability P_k, it becomes possible to calculate the expected value of perfect information [EVPI(A_i)], given that the decision maker would chose action A_i before receiving this information with the following expression:

$$\text{EVPI}(A_i) = \sum_k P_k R_{ik} \tag{1}$$

The above approach can be extended to the case of imperfect information (Raiffa 1968) by replacing P_k in the above equation with the probability of event k (E_k) given the imperfect sample information (I). This results in an expression for the expected value of sample information [EVSI(A_i, I)], given that the decision maker would chose action A_i before receiving this information:

$$\text{EVSI}(A_i, I) = \sum_k (P_k|I) R_{ik} \tag{2}$$

The value of imperfect (or sample) information provides a normative rule for deciding whether to collect additional information. For example, a decision to perform a survey before introducing a product can be made by comparing the cost of the survey to the expected value of the information obtained. It is often assumed that decision makers are biased when they fail to seek out additional information. The above discussion shows that *not* obtaining information is justified when the information costs too much. From a practical perspective, the value of information can guide decisions to provide information to product users (Lehto and Papastavrou 1991).

2.1.6. Maximizing Expected Value

From elementary probability theory, return is maximized by selecting the alternative with the greatest expected value. The expected value of an action A_i is calculated by weighting its consequences C_{ik} over all events k, by the probability P_{ik} the event will occur. The expected value of a given action A_i is therefore:

$$\text{EV}[A_i] = \sum_k P_{ik} C_{ik} \tag{3}$$

More generally, the decision maker's preference for a given consequence C_{ik} might be defined by a value function $V(C_{ik})$, which transforms consequences into preference values. The preference values are then weighted using the same equation. The expected value of a given action A_i becomes:

$$\text{EV}[A_i] = \sum_k P_{ik} V(C_{ik}) \tag{4}$$

Monetary value is a common value function. For example, lives lost, units sold, or air quality might all be converted into monetary values. More generally, however, value reflects preference, as illustrated by ordinary concepts such as the value of money or the attractiveness of a work setting. Given that the decision maker has large resources and is given repeated opportunities to make the choice, choices made on the basis of expected monetary value are intuitively justifiable. A large company might make nearly all of its decisions on the basis of expected monetary value. Insurance buying and many other rational forms of behavior can not, however, be justified on the basis of expected monetary value. Many years ago, it was already recognized that rational decision makers made choices not easily explained by expected monetary value (Bernoulli 1738). Bernoulli cited the St. Petersburg paradox, in which the prize received in a lottery was 2^n and n was the number of times (n) a flipped coin turned up heads before a tails was observed. The probability of n flips before the first tail is observed is 0.5^n. The expected value of this lottery becomes:

$$\text{EV}[L] = \sum_k P_{ik} V(C_{ik}) = \sum_{n=0}^{\infty} 0.5^n 2^n = \sum_{n=1}^{\infty} 1 \Rightarrow \infty \tag{5}$$

The interesting twist is that the expected value of the above lottery is infinite. Bernoulli's conclusion was that preference cannot be a linear function of monetary value, since a rational decision maker would never pay more than a finite amount to play the lottery. Furthermore, the value of the lottery can vary between decision makers. According to utility theory, this variability reflects rational differences in preference between decision makers for uncertain consequences.

2.1.7. Subjective Expected Utility (SEU) Theory

Expected utility theory extended expected value theory to describe better how people make uncertain economic choices (von Neumann and Morgenstern 1947). In their approach, monetary values are first transformed into utilities, using a utility function $u(x)$. The utilities of each outcome are then weighted by their probability of occurrence to obtain an expected utility. Subjective utility theory (SEU) added the notion that uncertainty about outcomes could be represented with subjective probabilities (Savage 1954). It was postulated that these subjective estimates could be combined with evidence using Bayes' rule to infer the probabilities of outcomes* (see Section 2.2). This group of assumptions corresponds to the Bayesian approach to statistics. Following this approach, the SEU of an alternative (A_i), given subjective probabilities (S_{ik}) and consequences (C_{ik}) over the events E_k, becomes:

$$\text{SEU}[A_i] = \sum_k S_{ik} U(C_{ik}) \tag{6}$$

Note the similarity between the above formulation for SEU and the earlier equation for expected value. EV and SEU are equivalent if the value function equals the utility function. Methods for eliciting value and utility functions differ in nature (Section 3). Preferences elicited for uncertain outcomes measure utility.† Preferences elicited for certain outcomes measure value. It accordingly has often been assumed that value functions differ from utility functions, but there are reasons to treat value and utility functions as equivalent (Winterfeldt and Edwards 1986). The latter authors claim that the differences between elicited value and utility functions are small and that "severe limitations constrain those relationships, and only a few possibilities exist, one of which is that they are the same."

When people are presented choices that have uncertain outcomes, they react in different ways. In some situations, people find gambling to be pleasurable. In others, people will pay money to reduce uncertainty; for example, when people buy insurance. SEU theory distinguishes between risk neutral, risk averse, risk seeking, and mixed forms of behavior. These different types of behavior are described by the shape of the utility function (Figure 2).

A risk-neutral decision maker will find the expected utility of a gamble to be the same as the utility of the gamble's expected value. That is, expected u(gamble) = u(gamble's expected value). For a risk-averse decision maker, expected u(gamble) < u(gamble's expected value); for a risk-

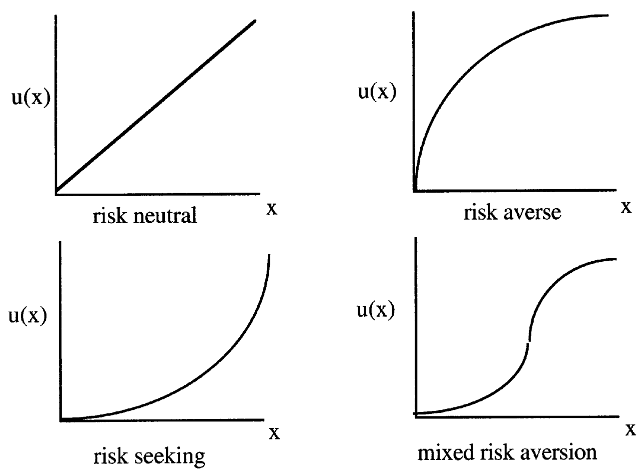

Figure 2 Utility Functions for Differing Risk Attitudes.

* When no evidence is available concerning the likelihood of different events, it was postulated that each consequence should be assumed to be equally likely. The Laplace decision rule makes this assumption and then compares alternatives on the basis of expected value or utility.

† Note that classical utility theory assumes that utilities are constant. Utilities may, of course, fluctuate. The random utility model (Bock and Jones 1968) allows such fluctuation.

seeking decision maker, expected u(gamble) > u(gamble's expected value). On any given point of a utility function, attitudes towards risk are described formally by the coefficient of risk aversion:

$$C_{RA} = \frac{u''(x)}{u'(x)} \tag{7}$$

where $u'(x)$ and $u''(x)$ are respectively the first and second derivatives of $u(x)$ taken with respect to x. Note that when $u(x)$ is a linear function of x, that is, $u(x) = ax + b$, then $C_{RA} = 0$. For any point of the utility function, if $C_{RA} < 0$, the utility function depicts risk-averse behavior, and if $C_{RA} > 0$, the utility function depicts risk-seeking behavior. The coefficient of risk aversion therefore describes attitudes toward risk at each point of the utility function, given that the utility function is continuous. SEU theory consequently provides a powerful tool for describing how people might react to uncertain or risky outcomes. However, some commonly observed preferences between risky alternatives can not be explained by SEU. Section 4.2 focuses on experimental findings showing deviations from the predictions of SEU.

A major contribution of SEU is that it represents differing attitudes towards risk and provides a normative model of decision making under uncertainty. The prescriptions of SEU are also clear and testable. Consequently, SEU has played a major role in fields other than economics, both as a tool for improving human decision making and as a stepping stone for developing models that describe how people make decisions when outcomes are uncertain. As discussed further in Section 4, much of this work has been done in psychology.

2.1.8. Multiattribute Utility Theory

Multiattribute tility theory (Keeney and Raiffa 1976) extends SEU to the case where the decision maker has multiple objectives. The approach is equally applicable for describing utility and value functions. Following this approach, the utility (or value) of an alternative A, with multiple attributes x, is described with the multiattribute utility (or value) function $u(x_1 \ldots x_n)$, where $u(x_1 \ldots x_n)$ is some function $f(x_1 \ldots x_n)$ of the attributes x. In the simplest case, multiattribute utility theory (MAUT) describes the utility of an alternative as an additive function of the single attribute utility functions $u_n(x_n)$. That is,

$$u(x_1 \ldots x_n) = \sum_n k_n u_n(x_n) \tag{8}$$

where the constants k_n are used to weight each single attribute utility function (u_n) in terms of its importance. Assuming an alternative has three attributes, x, y, and z, an additive utility function is $u(x,y,z) = k_x u_x(x) + k_y u_y(y) + k_z u_z(z)$. Along these lines, a community considering building a bridge across a river vs. building a tunnel or continuing to use the existing ferry system might consider the attractiveness of each option in terms of the attributes of economic benefits, social benefits, and environmental benefits.*

More complex multiattribute utility functions, include multiplicative forms and functions that combine utility functions for subsets of two or more attributes (Keeney and Raiffa 1976). An example of a simple multiplicative function would be $u(x,y) = u_x(x)*u_y(y)$. A function that combines utility functions for subsets, would be $u(x,y,z) = k_{xy}u_{xy}(x,y) + k_z u_z(z)$. This latter type of function becomes useful when utility independence is violated. Utility independence is violated when the utility function for one attribute depends on the value of another attribute. Along these lines, when assessing $u_{xy}(x,y)$, it might be found that $u_x(x)$ depends on the value of y. For example, peoples' reaction to the level of crime in their own neighborhood might depend on the level of crime in a nearby suburb. In the latter case, it is probably better to directly measure u_{xy}(x = crime in own neighborhood, y = crime in nearby suburb) than to estimate it from the single-attribute functions. The assessment of utility and value functions is discussed later in Section 3.

MAUT has been applied to a wide variety of problems (Saaty 1988; Keeney and Raiffa 1976; Winterfeldt and Edwards 1986; Clemen 1996). An advantage of MAUT is that it helps structure complex decisions in a meaningful way. Alternative choices and their attributes often naturally divide into hierarchies. The MAUT approach encourages such divide-and-conquer strategies and, especially in its additive form, provides a straightforward means of recombining weights into a final ranking of alternatives. The MAUT approach is also a compensatory strategy that allows normative trade-offs between attributes in terms of their importance.

*To develop the multiattribute utility function, the single-attribute utility functions (u_n) and the importance weights (k_n) are determined by assessing preferences between alternatives. Methods of doing so are discussed in Section 3.4.

2.1.9. Holistic Comparison

Holistic comparison is a nonanalytical method of comparing alternatives. This process involves a holistic comparison of the consequences for each alternative instead of separately measuring and then recombining measures of probability, value, or utility (Sage 1981; Stanoulov 1994; Janis and Mann 1977). A preference ordering between alternatives is thus obtained. For example, the decision maker might rank in order of preference a set of automobiles that vary on objectively measureable attributes, such as color, size, and price. Mathematical tools can then be used to derive the relationship between observed ordering and attribute values and ultimately predict preferences for unevaluated alternatives, as discussed in Section 3.3.4.

One advantage of holistic comparison is that it requires no formal consideration of probability or utility. Consequently, decision makers unfamiliar with these concepts may find holistic comparison to be more intuitive, and potential violations of the axioms underlying SEU and MAUT, due to their lack of understanding, become of lesser concern. People seem to find the holistic approach helpful when they compare complex alternatives (Janis and Mann 1977). In fact, people may feel there is little additional benefit to be obtained from separately analyzing the probability and value attached to each attribute. This tendency becomes prevalent in naturalistic decision making, as addressed further in Section 5.

2.2. Statistical Inference

Inference is the procedure followed when a decision maker uses information to determine whether a hypothesis about the world is true. Hypotheses can specify past, present, or future states of the world, or causal relationships between variables. Diagnosis is concerned with determining past and present states of the world. Prediction is concerned with determining future states. Inference or diagnosis is required in many decision contexts. For example, before deciding on a treatment, a physician must first diagnose the illness.

From the classical perspective, the decision maker is concerned with determining the likelihood that a hypothesis (H_i) is true. Bayesian inference is the best-known technique, but signal detection theory, and fundamentally different approaches such as the Dempster–Schafer method, have seen application. Each of these approaches is discussed below.

2.2.1. Bayesian Inference

Bayesian inference is a well-defined procedure for inferring the probability (P_i) that a hypothesis (H_i) is true, from evidence (E_j) linking the hypothesis to other observed states of the world. The approach makes use of Bayes' rule to combine the various sources of evidence (Savage 1954). Bayes' rule states that the posterior probability of hypothesis H_i given that evidence E_j is present, or $P(H_i|E_j)$, is given by the equation:

$$P(H_i|E_j) = \frac{P(E_j|H_i)P(H_i)}{P(E_j)} \tag{9}$$

where $P(H_i)$ is the probability of the hypothesis being true prior to obtaining the evidence E_j and $P(E_j|H_i)$ is the probability of obtaining the evidence E_j given that the hypothesis H_i is true. For example, consider the case where a physician is attempting to determine whether a patient has a disease present in 10% of the general population. The physician has a test available that gives a positive result 90% of the time when administered to patients who actually have the disease. The test also gives a positive result 20% of the time when administered to patients who don't have the disease. If the test were to be administered to a member of the general population, Eq. (9) predicts that the probability of having the disease given a positive test result is:

$$P(\text{disease}|\text{positive test}) = \frac{P(\text{positive test}|\text{disease})P(\text{disease in general population})}{P(\text{positive test})}$$

Also,

$$\begin{aligned} P(\text{positive test}) &= P(\text{positive test}|\text{disease})P(\text{disease in general population}) \\ &\quad + P(\text{positive test}|\text{no disease})P(\text{no disease in general population}) \\ P(\text{disease}|\text{positive test}) &= \frac{0.9*0.1}{0.9*0.1 + 0.2*0.9} = 0.33 \end{aligned}$$

As discussed further in Section 4.1, people often fail to combine evidence consistently with the above predictions of Bayes' rule. A common finding is that people fail to adequately consider the base rate of the hypothesis. In the above example, this would correspond to focusing on P(positive test|disease) = 0.9 and not considering P(disease in general population) = 0.1. As a consequence, many people might be surprised that P(disease|positive test) = 0.33 rather than a number close to 0.9.

When the evidence E_j consists of multiple states $E_1, \ldots, E_n$, each of which is conditionally independent, Bayes' rule can be expanded into the expression:

$$P(H_i|E_j) = \frac{\prod_{j=1}^{n} P(E_j|H_i)P(H_i)}{P(E_j)} \tag{10}$$

Calculating $P(E_j)$ can be somewhat difficult, due to the fact that each piece of evidence must be dependent,* or else it would not be related to the hypothesis. The odds forms of Bayes' rule provides a convenient way of looking at the evidence for and against a hypothesis that doesn't require $P(E_j)$ to be calculated. This results in the expression:

$$\Phi(H_i|E_j) = \frac{P(H_i|E_j)}{P(\sim H_i|E_j)} = \frac{\prod_{j=1}^{n} P(E_j|H_i)P(H_i)}{\prod_{j=1}^{n} P(E_j)|\sim H_i)P(\sim H_i)} \tag{11}$$

where $\Phi(H_i|E_j)$ refers to the posterior odds for hypothesis H_i, $P(\sim H_i)$ is the prior probability that hypothesis H_i is not true, and $P(\sim H_i|E_j)$ is the posterior probability that hypothesis H_i is not true.

The two latter forms of Bayes' rule provide an analytically simple way of combining multiple sources of evidence. Bayesian inference becomes much more difficult when the evidence is not certain or when the conditional independence assumption is not met. When evidence is not certain, complex multistage forms of Bayesian analysis are required that consider the probability of the evidence being true (Winterfeldt and Edwards 1986). When conditional independence is not true, the expanded form of Bayes' rule must be modified. For example, consider the case where the evidence consists of three events (E_1, E_2, E_3), where E_1 and E_2 are conditionally dependent and E_3 is conditionally independent of the two other events. The posterior probability, $P(H_i|E_1,E_2,E_3)$, then becomes:

$$P(H_i|E_1,E_2,E_3) = \frac{P(E_1,E_2|H_i)P(E_3|H_i)P(H_i)}{P(E_1,E_2)P(E_3|E_1,E_2)} \tag{12}$$

where $P(E_1,E_2|H_i)$ is the conditional probability of obtaining E_1 and E_2 given the hypothesis H_i, $P(E_3|H_i)$ is the conditional probability of obtaining E_3 given H_i, and $P(E_1,E_2)P(E_3|E_1,E_2)$ is the probability of obtaining the evidence (E_1, E_2, E_3).

2.2.2. Signal-Detection Theory

Bayesian inference combined with SEU leads to signal-detection theory (Tanner and Swets 1954), which has been applied in a large variety of contexts to model human performance (Wickens 1992). In signal-detection theory, the human operator is assumed to use Bayes' rule to estimate the probability that a signal actually is present from a noisy observation of the system. For example, an operator might estimate the probability a machine is going out of tolerance from a warning signal. The responses of the operator and the true state of the system together determine a set of four outcomes (Table 4).

TABLE 4 Potential Outcomes Considered by Signal Detection Theory

		State of the World	
		Noise (N)	Signal (S)
Response	yes	false alarm (fa)	hit (h)
	no	correct rejection (cr)	miss (m)

*Note that conditional independence between E_1 and E_2 implies that $P(E_1/H_i,E_2) = P(E_1/H_i)$ and that $P(E_2/H_i,E_1) = P(E_2/H_i)$. This is very different from simple independence, which implies that $P(E_1) = P(E_1/E_2)$ and that $P(E_2) = P(E_2/E_1)$.

The signal-detection model assumes an operator receives evidence from the environment regarding the true state of the world. The relationship between the signal (S) and the evidence (E) is measured by the conditional probability [$P(E|S)$] of obtaining the observed evidence given the signal is there. The decision maker is assumed to select a criterion value (x_c) that the evidence must exceed before saying yes. It is assumed that the value chosen will maximize utility. If the evidence is represented with a variable x, the expected utility of the operator can be described in terms of x, x_c and the four outcomes in Table 4. The expected utility for a given probability cutoff x_c, and utility function u, is given by the expression:

$$\begin{aligned}\mathrm{SEU}[x_c] &= P(x \geq x_c|S)P(S)u(h) + P(x \geq x_c|N)P(N)u(fa) \\ &\quad + P(x < x_c|S)P(S)u(m) + P(x < x_c|N)P(N)u(cr)\end{aligned} \tag{13}$$

where h is a hit, fa is a false alarm, m is a miss, and cr is a correct rejection. The above expression can be maximized by first substituting $1 - P(x \geq x_c|N)$ for $P(x < x_c|N)$ and also substituting $1 - P(x \geq x_c|S)$ for $P(x < x_c|S)$ into the equation for SEU[x_c], and then setting the derivative of SEU[x_c] with respect to x_c to zero. The result at the cutoff x_c is shown below:

$$\frac{P(x = x_c|S)}{P(x = x_c|N)} = \beta^* \geq \frac{P(N)(u(cr) - u(fa))}{P(S)(u(h) - u(m))} \tag{14}$$

where β^* is the optimal value of β. Substituting back the relation between $P(E|S)$ and the evidence x, the optimal decision rule is to say yes if

$$\frac{P(E|S)}{P(E|N)} \geq \beta^* \tag{15}$$

Equation (15) can be extended to multiple operators or multiple sources of evidence (Lehto and Papastavrou 1991). The resulting expression takes into account the probability of a false alarm and the probability of detection for the other source of information. Lehto and Papastavrou use this approach to analyze situations where the other source of information is a warning signal. The extent to which human judgments correspond to the predictions of Bayes' rule is further discussed in Section 4.1.

2.2.3. Dempster–Schafer Method

The Dempster–Schafer method (Schafer 1976; Fedrizzi et al. 1994) is an alternative to Bayesian inference for accumulating evidence for or against a hypothesis that has been proposed for use in decision analysis (Strat 1994). In this approach, the relation of hypotheses (H) to evidence (e) is described by a basic probability assignment (bpa) function, p. Given evidence (e), this function $p_e(n)$ assigns a value between 0 and 1 to each subset of H, such that the sum of the values assigned is 1. For example, consider the case where there are three hypotheses (A,B,C). When no evidence is available, the vacuous bpa assigns a value of 1 to the set of hypotheses $H = (A,B,C)$ and a 0 to all subsets. That is, the subsets (A), (B), (C), (A,B), and (A,C) are each assigned a value of 0. The Bayesian approach would instead assign a probability of 0.33 to A, B, and C respectively.

Also, given that evidence $p_e(A) = x$ supporting a specific hypothesis A is found, the Dempster–Schafer approach assigns $(1 - p_e(A))$ to H. The Bayesian approach, of course, assigns $(1 - p_e(A))$ to the complement of A. Returning to the above example, suppose the evidence supports hypothesis A to the degree $p_e(A) = 0.6$. Using the Dempster–Schafer approach, $p_e(A,B,C) = 0.4$. This, of course, is very different from the Bayes' interpretation, where $P(A) = 0.6$ and $P(\text{not } A) = 0.4$. The Dempster–Schafer method uses a belief function $B(n)$ to assign a total belief to n, where n is a subset of the set of possible hypotheses (H), as the sum of the beliefs assigned to m, where m is the set of possible subsets of n. In the above example, the belief in (A,B,C) after receiving evidence (e) is as given below:

$$\begin{aligned}B(A,B,C) &= p_e(A,B,C) + p_e(A,B) + \mathrm{p}_e(A,C) + p_e(B,C) + p_e(A) + p_e(B) + p_e(C) \\ &= 0.4 + 0 + 0 + 0 + 0.6 + 0 + 0 \\ &= 1.0\end{aligned} \tag{16}$$

Similarly, the belief in (A,B) after receiving the evidence (e) is:

TABLE 5 Tableau for Dempster–Shafer Method of Combining Evidence

X = [(A), (A,B,C)]	Y = [(A,B), (A,B,C)] $p_f(A,B) = 0.4$	$p_f(A,B,C) = 0.6$
$p_e(A) = 0.6$	A ; $p_e p_f = 0.24$	A ; $p_e p_f = 0.36$
$p_e(A,B,C) = 0.4$	A,B ; $p_e p_f = 0.16$	A,B,C ; $p_e p_f = 0.24$

$$\begin{aligned} B(A,B) &= p_e(A,B) + p_e(A) + p_e(B) \\ &= 0 + 0.6 + 0 \\ &= 0.6 \\ &= B(A) \end{aligned} \tag{17}$$

To combine evidence from multiple sources e and f, Dempster–Schafer theory uses the combining function $c(p_e(X),p_f(Y))$, where X and Y are both sets of subsets of H. For example, we might have $X = [(A), (A,B,C)]$ and $Y = [(A,B), (A,B,C)]$. The combining function then assigns a value to each subset n of H. The value assigned is determined by first describing the set of subsets n' within n defined by the intersection of subsets within X and subsets within Y. A value of 0 is assigned to all subsets of n not within n'. The products $p_e(X)*p_f(Y)$ are then summed and assigned to each subset within n'. Returning to the above example, we can calculate $c(n')$ using the values given in Table 5. First note that the set of subsets n' for the example is defined by the inner elements of the table. Specifically, $n' = [(A), (A,B), (A,B,C)]$. The values used by the combining function $c(n')$ are also shown. Using these numbers, the values of $c(n')$ become: $c(A) = 0.24 + 0.36 = 0.6$; $c(A,B) = 0.16$; $c(A,B,C) = 0.24$. All remaining subsets for this evidence are assigned a value of 0.

It has been argued that the Dempster–Schafer method of assigning evidence is better suited for diagnosing medical problems than the Bayesian method (Gordon and Shortliffe 1984). The latter researchers particularly criticize the Bayesian assumption that evidence partially supporting a hypothesis should also support its negation. Gordon and Shortliffe note that the Dempster–Schafer method shows promise as a means of accumulating belief in expert diagnostic systems used in medicine.

3. DECISION ANALYSIS

The application of classical decision theory to improve human decision making is the goal of decision analysis (Raiffa 1968; Howard 1968; 1988; Keeney and Raiffa 1976). Decision analysis requires inputs from decision makers, such as goals, preference and importance measures, and subjective probabilities. Elicitation techniques have consequently been developed that help decision makers provide these inputs. Particular focus has been placed on methods of quantifying preferences, trade-offs between conflicting objectives, and uncertainty (Keeney and Raiffa 1976; Raiffa 1968). As a first step in decision analysis, it is necessary to do some preliminary structuring of the decision, which then guides the elicitation process. The following discussion first presents methods of structuring decisions and then covers techniques for assessing subjective probabilities, utility functions, and preferences.

3.1. Structuring Decisions

The field of decision analysis has developed many useful frameworks for representing what is known about a decision (Howard 1968; Winterfelt and Edwards 1986; Clemen 1996). In fact, the above authors and others have stated that the process of structuring decisions is often the greatest contribution of going through the process of decision analysis. Among the many tools used, decision matrices and trees provide a convenient framework for comparing decisions on the basis of expected value or utility. Value trees provide a helpful method of structuring the sometimes complex relationships among objectives, attributes, goals, and values and are used extensively in multiattribute decision-making problems. Event trees, fault trees, inference trees, and influence diagrams are useful for describing probabilistic relationships between events and decisions. Each of these approaches is briefly discussed below.

3.1.1. Decision Matrices and Trees

Decision matrices are often used to represent single-stage decisions (Figure 3). The simplicity of decision matrices is their primary advantage. They also provide a very convenient format for applying

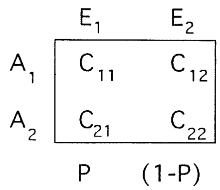

Figure 3 Decision Matrix Representation of a Single-Stage Decision.

the decision rules discussed in the previous section. Decision trees are also commonly used to represent single-stage decisions (Figure 4) and are particularly useful for describing multistage decisions (Raiffa 1968). Note that in a multistage decision tree, the probabilities of later events are conditioned on the result of earlier events. This leads to the important insight that the results of earlier events provide information regarding future events.* Following this approach, decisions may be stated in conditional form. An optimal decision, for example, might be to first do a market survey, and then market the product only if the survey to is positive.

Analysis of a single or multistage decision tree involves two basic steps referred to as averaging out and folding back (Raiffa 1968). These steps, respectively, occur at chance and decision nodes.† Averaging out occurs when the expected value (or utility) at each chance node is calculated. In Figure 4, this corresponds to calculating the expected value of A_1 and A_2, respectively. Folding back refers to choosing the action with the greatest expected value at each decision node.

Decision trees consequently provide a straightforward way of comparing alternatives in terms of expected value or SEU. However, their development requires significant simplification of most decisions and the provision of numbers, such as measures of preference and subjective probabilities, that decision makers may have difficulty determining. In certain contexts, decision makers struggling with this issue may find it helpful to develop value trees, event trees, or influence diagrams, as expanded upon below.

3.1.2. Value Trees

Value trees hierarchically organize objectives, attributes, goals, and values (Figure 5). From this perspective, an objective corresponds to satisficing or maximizing a goal or set of goals. When there is more than one goal, the decision maker will have multiple objectives, which may differ in importance. Objectives and goals are both measured on a set of attributes. Attributes may provide (1) objective measures of an goal, such as when fatalities and injuries are used as a measure of highway safety, (2) subjective measures of an goal, such as when people are asked to rate the quality of life in the suburbs vs. the city, or (3) proxy or indirect measures of a goal, such as when the quality of ambulance service is measured in terms of response time.

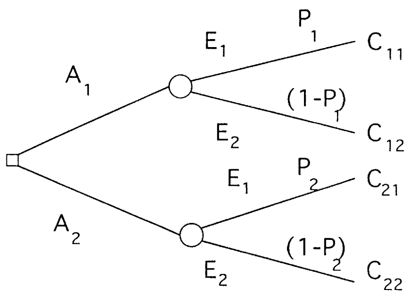

Figure 4 Decision Tree Representation of a Single-Stage Decision.

* For example, the first event in a decision tree might be the result of a test. The test result then provides information useful in making the final decision.

† Note that standard convention uses circles to denote chance nodes and squares to denote decision nodes (Raiffa 1968).

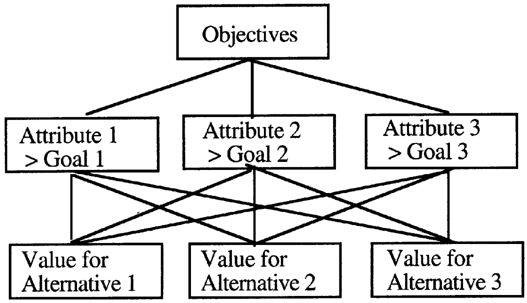

Figure 5 Generic Value Tree.

In generating objectives and attributes, it becomes important to consider their relevance, completeness, and independence. Desirable properties of attributes (Keeney and Raiffa 1976) include:

1. *Completeness:* The extent to which the attributes measure whether an objective is met.
2. *Operationality:* The degree to which the attributes are meaningful and feasible to measure.
3. *Decomposability:* Whether the whole is described by its parts.
4. *Nonredundancy:* Correlated attributes give misleading results.
5. *Minimum size:* Considering irrelevant attributes is expensive and may be misleading.

Once a value tree has been generated, various methods can be used to assess preferences directly between the alternatives.

3.1.3. Event Trees or Networks

Event trees or networks show how a sequence of events can lead from primary events to one or more outcomes. Human reliability analysis (HRA) event trees are a classic example of this approach (Figure 6). If probabilities are attached to the primary events, it becomes possible to calculate the probability of outcomes, as illustrated in Section 3.2.4. This approach has been used in the field of risk assessment to estimate the reliability of human operators and other elements of complex systems (Gertman and Blackman 1994). Chapter 32 provides additional information on human reliability analysis and other methods of risk assessment.

Fault trees work backwards from a single undesired event to its causes (Figure 7). Fault trees are commonly used in risk assessment to help infer the chance of an accident occurring (Hammer 1993; Gertman and Blackman 1994). Inference trees relate a set of hypotheses at the top level of the tree to evidence depicted at the lower levels. This latter approach has been used by expert systems, such as PROSPECTOR (Duda et al. 1979). PROSPECTOR applies a Bayesian approach to infer the presence of a mineral deposit from uncertain evidence.

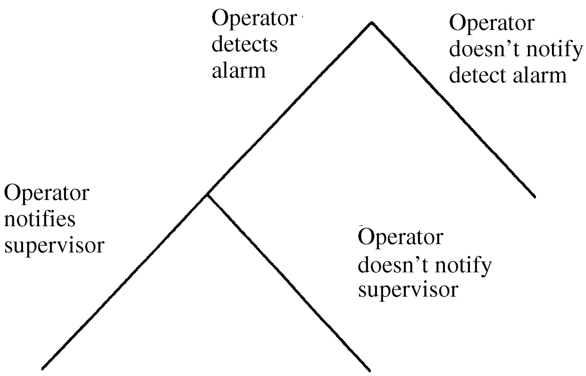

Figure 6 HRA Event Tree. (Adapted from Gertman and Blackman 1994)

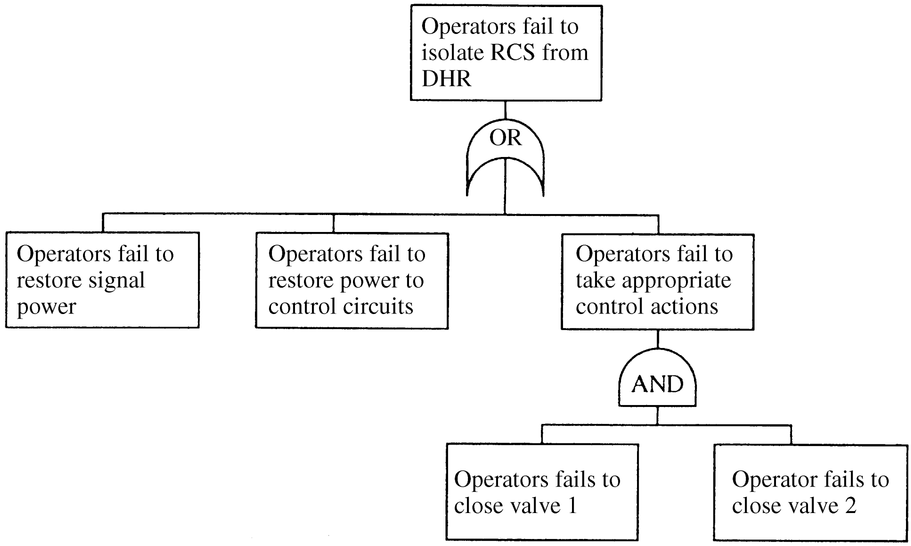

Figure 7 Fault Tree for Operators. (Adapted from Gertman and Blackman 1994).

3.1.4. Influence Diagrams and Cognitive Mapping

Influence diagrams are often used in the early stages of a decision to show how events and actions are related. Their use in the early stages of a decision is referred to as knowledge (or cognitive) mapping (Howard 1988). Links in an inference diagram depict causal and temporal relations between events and decision stages.* A link leading from an event *A* to an event *B* implies that the probability of obtaining event *B* depends on whether event *A* has occurred. A link leading from a decision to a event implies that the probability of the event depends on the choice made at that decision stage. A link leading from an event to a decision implies that the decision maker knows the outcome of the event at the time the decision is made.

One advantage of influence diagrams in comparison to decision trees is that influence diagrams show the relationships between events more explicitly. Consequently, influence diagrams are often used to represent complicated decisions where events interactively influence the outcomes. For example, the influence diagram in Figure 8 shows that the true state of the machine affects both the

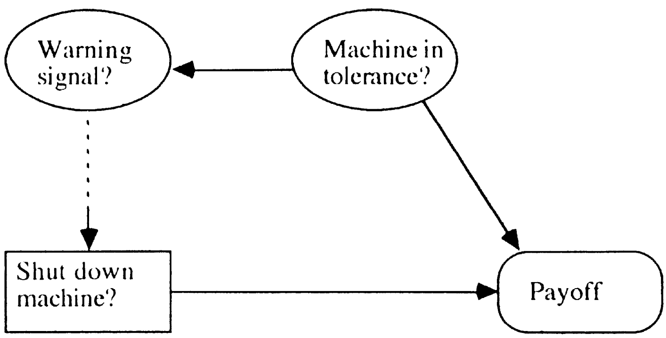

Figure 8 Influence Diagram Representation of a Single-Stage Decision.

*As for decision trees, the convention for influence diagrams is to depict events with circles and decisions with squares.

probability of the warning signal and the consequence of the operator's decision. This linkage would be hidden within a decision tree.* Influence diagrams have been used to structure medical decision-making problems (Holtzman 1989) and are emphasized in modern texts on decision analysis (Clemen 1996). Howard (1988) states that influence diagrams are the greatest advance he has seen in the communication, elicitation, and detailed representation of human knowledge. Part of the issue is that influence diagrams allow people who do not have deep knowledge of probability to describe complex conditional relationships with simple linkages between events. Once these linkages are defined, the decision becomes well defined and can be formally analyzed.

3.2. Probability Assessment

Several approaches have been used in decision analysis to assess subjective probabilities. In this section several of the more well-known techniques will be summarized. These techniques include: (1) direct numerical assessment, (2) fitting subjective belief forms, (3) the bisection method, (4) conditioning arguments, (5) preferences between reference gambles, and (6) scaling methods. Techniques proposed for improving the accuracy of assessed probabilities, including scoring rules, calibration, and group assessment, will then be presented.

3.2.1. Direct Numerical Assessment

In direct numerical estimation, decision makers are asked to give a numerical estimate of how likely they think the event is to happen. These estimates can be probabilities, odds, log odds, or words (Winterfeldt and Edwards 1986). Winterfeldt and Edwards argue that log odds have certain advantages over the other measures. Gertman and Blackman (1994) note that log odds are normally used in risk assessment for nuclear power applications because human error probabilities (HEPs) vary greatly in value. HEPs between 1 and 0.00001 are typical.

3.2.2. Fitting a Subjective Belief Form

Fitting a subjective belief form requires that the questions be posed in terms of statistical parameters. That is, decision makers could be asked to first consider their uncertainty regarding the true value of a given probability and then estimate their mean, mode, or median belief. This approach can be further extended by asking decision makers to describe how certain they are of their estimate. For example, a worker might subjectively estimate the mean and variance of the proportion of defective circuit boards before inspecting a small sample of circuit boards. If the best estimate corresponds to a mean, mode, or median, and the estimate of certainty to a confidence interval or standard deviation, a functional form such as the Beta-1 probability density function (pdf) can then be used to fit a subjective probability distribution (Clemen 1996; Buck 1989).

In other words, a distribution is specified that describes the subject's belief that the true probability equals particular values. This type of distribution can be said to express uncertainty about uncertainty (Raiffa 1968). Given that the subject's belief can be described with a Beta-1 pdf, Bayesian methods can be used to combine binomially distributed evidence easily with the subject's prior belief (Clemen 1996; Buck 1989). Returning to the above example, the worker's prior subjective belief can be combined with the results of inspecting the small sample of circuit boards, using Bayes' rule. As more evidence is collected, the weight given to the subject's initial belief becomes smaller compared to the evidence collected. The use of prior belief forms also reduces the amount of sample information that must be collected to show that a proportion, such as the percentage of defective items, has changed (Buck 1989).

3.2.3. Bisection Method

The bisection method (Raiffa 1968) is another direct technique for attempting to estimate a subjective probability density function (pdf). This technique is somewhat more general than fitting the subject's belief with a functional form, such as the beta-1, since it makes no parametric assumptions. The bisection method involves two steps which are repeated until the subject's belief is adequately described. Following this approach, the first step is to determine the median ($P_{0.5}$) of the subjective pdf. This question is posed to the decision maker in a form such as "For what value of p do you feel it is equally likely the true value p† is greater than or less than p?" This step is then repeated for subintervals to obtain the desired level of detail.

* The conditional probabilities in a decision tree would reflect this linkage, but the structure of the tree itself does not show the linkage directly. Also, the decision tree would use the flipped probability tree using P(warning) at the first stage and P(machine down|warning) at the second stage. It seems more natural for operators to think about the problem in terms of P(machine down) and P(warning|machine down), which is the way the influence diagram in Figure 7 depicts the relationship.

† Note that it has been shown that people viewing fault trees can be insensitive to missing information (Fischhoff et al. 1978).

3.2.4. Conditioning Arguments

Statistical conditioning arguments are based on the idea that the probability of a complicated event, such as the chance of having an accident, can be determined by estimating the probability of simpler events (or subsets). From a more formal perspective, a conditioning argument determines the probability of an event A by considering the possible conditions (C_i) under which A might happen, the associated conditional probabilities [$P(A|C_i)$], and the probability of each condition [$P(C_i)$]. The probability of A can then be represented as:

$$P(A) = \sum_i P(A|C_i)P(C_i) \tag{18}$$

This approach is illustrated by the development of event trees and fault tree analysis. In fault tree analysis, the probability of an accident is estimated by considering the probability of human errors, component failures, and other events. This approach has been extensively applied in the field of risk analysis (Gertman and Blackman 1994).* THERP (Swain and Guttman 1983) extends the conditioning approach to the evaluation of human reliability in complex systems.

SLIM-MAUD (Embrey 1984) implements a related approach in which expert ratings are used to estimate human error probabilities (HEPs) in various environments. The experts first rate a set of tasks in terms of performance-shaping factors (PSFs) that are present. Tasks with known HEPs are used as upper and lower anchor values. The experts also judge the importance of individual PSFs. A subjective likelihood index (SLI) is then calculated for each task in terms of the PSFs. A logarithmic relationship is assumed between the HEP and SLI, allowing calculation of the human error probability for task j (HEP_j) from the subjective likelihood index assigned to task j (SLI_j). More specifically:

$$\mathrm{Log}(1 - \mathrm{HEP}_j) = a\mathrm{SLI}_j + b \tag{19}$$

$$\text{where } \mathrm{SLI}_j = \sum_i \mathrm{PSF}_{ij} * I(\mathrm{PSF}_i) \tag{20}$$

$I(\mathrm{PSF}_i)$ is the importance of PSF_i, and PSF_{ij} is the rating given to PSF_i for task j. Gertman and Blackman (1994) provide guidelines regarding the use of this method and have generally positive conclusions. SLIM-MAUD is interesting in that it uses multiattribute utility theory as a basis for generating probability estimates.

3.2.5. Reference Lotteries

Reference lottery methods take a less direct approach to obtaining point estimates of the decision maker's subjective probabilities. When the objective is to measure how likely event A is to occur, the approach asks decision makers to consider a lottery where they will receive a prize x if event A occurs, and a prize y if it does not. They are then asked how much they would be willing to pay for the lottery. The amount they are willing to pay z is then equated to the lottery, using the relation $z = P(A)x + [1 - P(A)]y$. From this expression it becomes possible to estimate the decision maker's subjective estimate of $P(A)$. Specifically, $P(A) = (z - y)/(x - y)$. A variant of this approach that asks decision makers to compare two lotteries over the same range of preferences might be preferable because it removes the potential effect of risk aversion (Winterfelt and Edwards 1986).

3.2.6. Scaling Methods

Scaling methods ask subjects to rate or rank the probabilities to be assessed. Likert scales with verbal anchors have been used to obtain estimates of how likely people feel certain risks are (Kraus and Slovic 1988). Another approach has been to ask subjects to do pair-wise comparisons of the likelihoods of alternative events (Saaty 1988). Pairwise comparisons of probabilities on a ratio scale correspond to relative odds and consequently have high construct validity. In fact, much of the risk assessment focuses on determining order of magnitude differences in probability. Saaty (1988), however, argues that the psychometric literature indicates that people's ability to distinguish items on the same scale is limited to 7 ± 2 categories. He consequently proposes use of a relative scale to measure differences in importance, preference, and probability that uses verbal anchors corresponding to equal, weak, strong, very strong, and absolute differences between rated items. In perhaps the most controversial aspect of his approach, these five verbal anchors are assigned the numbers 1, 3, 5, 7, and 9. Using these numbers, subjective probabilities can then be calculated from pair-wise ratings on his verbal scale.

3.2.7. Scoring Rules, Calibration, and Group Assessment

A number of approaches have been developed for improving the accuracy of assessed probabilities (Winterfelt and Edwards 1986; Lichtenstein et al. 1982). Two desirable properties of elicited prob-

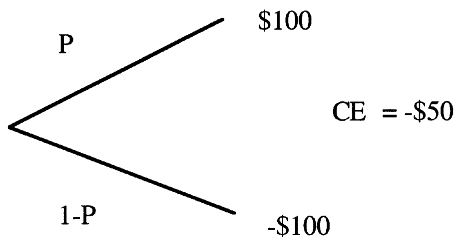

Figure 9 The Standard Gamble Used in the Variable Probability Method of Eliciting Utility Functions.

abilities include extremeness and calibration. More extreme probabilities (for example, P(good sales) = 0.9 vs. P(good sales) = 0.5) make decisions easier since the decision maker can be more sure of what is really going to happen. Well-calibrated probability estimates match the actual frequencies of observed events. Scoring rules provide a means of evaluating assessed probabilities in terms of both extremeness and calibration. If decision makers assess probabilities on a routine basis, feedback can be provided using scoring rules. Such feedback seems to be associated with the highly calibrated subjective probabilities provided by weather forecasters (Murphy and Winkler 1974).

Group assessment of subjective probabilities is another often-followed approach, as alluded to earlier in reference to SLIM-MAUD. There is evidence that group judgments are usually more accurate than individual judgments and that groups tend to be more confident in their estimates (Sniezek and Henry 1989; Sniezek 1992; also see Section 6.2). Assuming that individuals within a group independently provide estimates, which are then averaged, the benefit of group judgment is easily shown to have a mathematical basis. Simply put, a mean should be more reliable than an individual observation. Group dynamics, however, can lead to a tendency towards conformity (Janis 1972). Winterfeldt and Edwards (1986) therefore recommend that members of a group be polled independently.

3.3. Utility Function Assessment

Standard methods for assessing utility functions (Raiffa 1968) include (1) the variable probability method and (2) the certainty equivalent method. In the variable probability method, the decision maker is asked to give the value for the probability of winning at which they are indifferent between a gamble and a certain outcome (Figure 9). A utility function is then mapped out when the value of the certainty equivalent (CE) is changed over the range of outcomes. Returning to Figure 9, the value of P at which the decision maker is indifferent between the gamble and the certain loss of \$50 gives the value for $u(-\$50)$. In the utility function in Figure 10, the decision maker gave a value of about 0.5 in response to this question.

The certainty equivalent method uses lotteries in a similar way. The major change is that the probability of winning or losing the lottery is held constant, while the amount won or lost is changed. In most cases, the lottery provides an equal chance of winning and losing. The method begins by

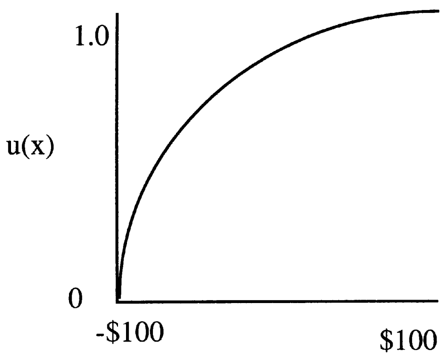

Figure 10 A Typical Utility Function.

asking the decision maker to give a certainty equivalent for the original lottery (CE_1). The value chosen has a utility of 0.5. This follows since the utility of the best outcome is assigned a value of 1 and the worst is given a utility of 0. The utility of the original gamble is therefore:

$$u(CE_1) = pu(\text{best}) + (1 - p)u(\text{worst}) = p(1) + (1 - p)(0) = p = 0.5 \tag{21}$$

The decision maker is then asked to give certainty equivalents for two new lotteries. Each uses the CE from the previous lottery as one of the potential prizes. The other prizes used in the two lotteries are the best and worst outcomes from the original lottery, respectively. The utility of the certainty equivalent (CE_2) for the lottery using the best outcome and CE_1 is given by the expression below:

$$u(CE_2) = pu(\text{best}) + (1 - p)u(CE_1) = p(1) + (1 - p)(0.5) = 0.75 \tag{22}$$

The utility of the certainty equivalent (CE_3) given for the lottery using the worst outcome and CE_1 is given by:

$$u(CE_3) = pu(CE_1) + (1 - p)u(\text{worst}) = p(0.5) + (1 - p)(0) = 0.25 \tag{23}$$

This process is continued until the utility function is specified in sufficient detail. A problem with the certainty equivalent method is that errors are compounded as the analysis proceeds. This follows since the utility assigned in the first preference assessment (i.e., $u(CE_1)$) is used throughout the subsequent preference assessments. A second issue is that the CE method uses different ranges in the indifference lotteries, meaning that the CEs are compared against different reference values. This might create inconsistencies since, as discussed later in Section 4, attitudes toward risk usually change depending upon whether outcomes are viewed as gains or losses. The use of different reference points may, of course, cause the same outcome to be viewed as either a loss or a gain. Utilities may also vary over time. Section 4.2 discusses some of these issues further.

3.4. Preference Assessment

Methods for measuring strength of preference include indifference methods, direct assessment, and indirect measurement (Keeney and Raiffa 1976; Winterfeldt and Edwards 1986). Indifference methods modify one of two sets of stimuli until subjects feel they are indifferent between the two. Direct-assessment methods ask subjects to rate or otherwise assign numerical values to attributes, which are then used to obtain preferences for alternatives. Indirect-measurement techniques avoid decomposition and simply ask for preference orderings between alternatives. Recently there has been some movement towards evaluating the effectiveness of particular methods for measuring preferences (Huber et al. 1993; Birnbaum et al. 1992).

3.4.1. Indifference Methods

Indifference methods are illustrated by the variable probability and certainty equivalent methods of eliciting utility functions presented in the previous section. There, indifference points were obtained by varying either probabilities or values of outcomes. Similar approaches have been applied to develop multiattribute utility or value functions. This approach involves four steps: (1) develop the single attribute utility or value functions, (2) assume a functional form for the multiattribute function, (3) assess the indifference point between various multiattribute alternatives, and (4) calculate the substitution rate or relative importance of one attribute compared to the other. The single-attribute functions might be developed by indifference methods (i.e., the variable probability or certainty equivalent methods) or direct-assessment methods, as discussed later. Indifference points between multiattribute outcomes are obtained through an interactive process in which the values of attributes are systematically increased or decreased. Substitution rates are then obtained from the indifference points.

For example, consider the case for two alternative traffic safety policies, A_1 and A_2. Each policy has two attributes, x = lives lost and y = money spent. Assume the decision maker is indifferent between A_1 and A_2, meaning the decision maker feels that $v(x_1,y_1) = v(20{,}000 \text{ deaths; } \$1 \text{ trillion})$ is equivalent to $v(x_2,y_2) = v(10{,}000 \text{ deaths; } \$1.5 \text{ T})$. For the sake of simplicity, assume an additive value function, where $v(x,y) = (1 - k)v_x(x) + kv_y(y)$. Given this functional form, the indifference point $A_1 = A_2$ is used to derive the relation:

$$(1 - k)v_x(20{,}000 \text{ deaths}) + kv_y(\$1 \text{ T}) = (1 - k)v_x(10{,}000 \text{ deaths}) + kv_y(\$1.5 \text{ T}) \tag{24}$$

This results in the substitution rate as shown below:

$$\frac{k}{1-k} = \frac{v_x(20{,}000 \text{ deaths}) - v_x(10{,}000 \text{ deaths})}{v_y(\$1.5 \text{ T}) - v_y(\$1 \text{ T})} \tag{25}$$

If $v_x = -x$, and $v_y = -y$, a value of approximately 2^{-5} is obtained for k. The procedure becomes somewhat more complex when nonadditive forms are assumed for the multiattribute function (Keeney and Raiffa 1976).

3.4.2. *Direct-Assessment Methods*

Direct-assessment methods include curve fitting and various numerical rating methods (Winterfeldt and Edwards 1986). Curve fitting is perhaps the simplest approach. Here, the decision maker first orders the various attributes and then simply draws a curve assigning values to them. For example, an expert might draw a curve relating levels of traffic noise (measured in decibels) to their level of annoyance (on a scale of 0 to 1). Rating methods, as discussed earlier in reference to subjective probability assessment, include direct numerical measures on rating scales and relative ratings.

The analytic hierarchy process (AHP) provides one of the more implementable methods of this type (Saaty 1988). In this approach, the decision is first structured as a value tree (Figure 5). Then each of the attributes is compared in terms of importance in a pair-wise rating process. When entering the ratings, decision makers can enter numerical ratios (for example, an attribute might be twice as important as another) or use the subjective verbal anchors mentioned earlier in reference to subjective probability assessment. The AHP program uses the ratings to calculate a normalized eigenvector assigning importance or preference weights to each attribute. Each alternative is then compared on the separate attributes. For example, two houses might first be compared in terms of cost and then be compared in terms of attractiveness. This results in another eigenvector describing how well each alternative satisfies each attribute. These two sets of eigenvectors are then combined into a single vector that orders alternatives in terms of preference. The subjective multiattribute rating technique (SMART) developed by Edwards (see Winterfeldt and Edwards 1986) provides a similar, easily implemented approach. Both techniques are computerized, making the assessment process relatively painless.

3.4.3. *Indirect Measurement*

Indirect-measurement techniques avoid asking people to rate or rank directly the importance of factors that impact their preferences. Instead, subjects simply state or order their preferences for different alternatives. A variety of approaches can then be used to determine how individual factors influence preference. Conjoint measurement theory provides one such approach for separating the effects of multiple factors when only their joint effects are known. Application of the approach entails asking subjects to develop an ordered set of preferences for a set of alternatives that systematically vary attributes felt to be related to preference. The relationship between preferences and values of the attributes is then assumed to follow some functional form. The most common functional form assumed is a simple additive-weighting model. Preference orderings obtained using the model are then compared to the original rankings. Example applications of conjoint measurement theory to describe preferences between multiattribute alternatives are discussed in Winterfeldt and Edwards (1986). Related applications include the dichotomy-cut method, used to obtain decision rules for individuals and groups from ordinal rankings of multiattribute alternatives (Stanoulov 1994).

The policy capturing approach used in social judgment theory (Hammond et al. 1975; Hammond 1993) is another indirect approach for describing human judgments of both preferences and probability. The policy-capturing approach uses multivariate regression or other similar techniques to relate preferences to attributes for one or more decision makers. The obtained equations correspond to policies followed by particular decision makers. An example equation might relate medical symptoms to a physician's diagnosis. It has been argued that the policy-capturing approach measures the influence of factors on human judgments more accurately than decomposition methods. Captured weights might be more accurate because decision makers may have little insight into the factors that impact their judgments (Valenzi and Andrews 1973). People may also weigh certain factors in ways that reflect social desirability rather than influence on their judgments (Brookhouse et al. 1986). For example, people comparing jobs might rate pay as being lower in importance than intellectual challenge, while their preferences between jobs might be predicted entirely by pay. Caution must also be taken when interpreting regression weights as indicating importance, since regression coefficients are influenced by correlations between factors, their variability, and their validity (Stevenson et al. 1993).

4. BEHAVIORAL DECISION THEORY

As a normative ideal, classical decision theory has influenced the study of decision making in a major way. Much of the earlier work in behavioral decision theory compared human behavior to the prescriptions of classical decision theory (Edwards 1954; Slovic et al. 1977; Einhorn and Hogarth 1981). Numerous departures were found, including the influential finding that people use heuristics during judgment tasks (Tversky and Kahneman 1974). On the basis of such research, pyschologists have concluded that other approaches are needed to describe the process of human decision making. Descriptive models that relax assumptions of the normative models, but still retain much of their essence, are now being evaluated in the field of judgment and decision theory (Stevenson et al. 1993).

The following discussion summarizes findings from this broad body of literature. The discussion begins by considering research on statistical estimation and inference. Attention then shifts to the topic of decision making under uncertainty and risk.

4.1. Statistical Estimation and Inference

The ability of people to perceive, learn, and draw inferences accurately from uncertain sources of information has been a topic of much research. The following discussion first briefly considers human abilities and limitations on such tasks. The next section introduces several heuristics people seem to use to cope with their limitations and considers how their use can cause certain biases. Attention then shifts to probabilistic information-processing models and policy capturing models. These modeling approaches provide a mathematically oriented view of how people judge probabilities, the biases that might occur, and how people learn to perform probability judgment tasks. The final section briefly summarizes findings on debiasing human judgments.

4.1.1. Human Abilities/Limitations

Research conducted in the early 1960s tested the notion that people behave as ''intuitive statisticians'' who gather evidence and apply it in accordance with the Bayesian model of inference (Peterson and Beach 1967). Several studies evaluated how good people are at estimating statistical parameters, such as means, variances, and proportions. Other studies have compared human inferences obtained from probabilistic evidence to the prescriptions of Bayes' rule.

A number of interesting results were obtained (Table 6). The research first shows that people can be fairly good at estimating means, variances, or proportions from sample data. However, this ability drops greatly when the judged events occur either rarely or very often. In particular, when people are asked to estimate the risk associated with the use of consumer products (Dorris and Tabrizi 1978; Rethans 1980) or various technologies (Lichtenstein et al. 1978), estimates can be weakly related to accident data. Weather forecasters are one of the few groups of people that have been documented as being able to estimate high and low probabilities accurately (Winkler and Murphy 1973).

Part of the issue is that risk estimates are related to factors other than likelihood, such as catastrophic potential, degree of control, or familiarity (Lichtenstein et al. 1978; Slovic 1978; 1987; Lehto et al. 1994). Weber (1994) provides additional evidence that subjective probabilities are related to factors other than uncertainty and argues that people will overestimate the chance of highly positive outcome because of their desire to obtain it. Weber also argues that people will overestimate the chance of a highly undesirable outcome because of their fear of receiving it. Traditional methods of decision analysis separately elicit and then recombine subjective probabilities with utilities, as discussed earlier, and assume that subjective probabilities are independent of consequences. A finding of dependency therefore casts serious doubt upon the normative validity of this commonly accepted approach.

When studies of human inference are considered, several other trends become apparent (Table 6). In particular, several significant deviations from the Bayesian model have been found. These include:

1. Decision makers tend to be conservative in that they don't give as much weight to probabilistic evidence as Bayes' rule (Edwards 1968).
2. They don't consider base rates or prior probabilities adequately (Tversky and Kahneman 1974).
3. They tend to ignore the reliability of the evidence (Tversky and Kahneman 1974).
4. They tend to overestimate the probability of conjunctive events and underestimate the probability of disjunctive events (Bar-Hillel 1973).
5. They tend to seek out confirming evidence rather than disconfirming evidence and place more emphasis on confirming evidence when it is available (Einhorn and Hogarth 1978; Baron 1985).
6. They are overconfident in their predictions (Fischhoff et al. 1977), especially in hindsight (Fischhoff 1982; Christensen-Szalanski 1991).
7. They show a tendency to infer illusionary causal relations (Tversky and Kahneman 1973).

A lively literature has developed regarding these deviations and their significance (Evans 1989; Wickens 1992; Caverni et al. 1990; Klein et al. 1993). From one perspective, these deviations demonstrate inadequacies of human reason and are a source of societal problems (Hammond 1974). From the opposite perspective, it has been held that the above findings are more or less experimental

TABLE 6 Sample Findings on the Ability of People to Estimate and Infer Statistical Quantities

Statistical Estimation	
Accurate estimation of sample means	Peterson and Beach 1968
Variance estimates correlated with mean	Lathrop 1967
Variance biases not found	Levin 1975
Variance estimates based on range	Pitz 1980
Accurate estimates of sample proportions between 0.75 and 0.25	Edwards 1954
Severe overestimates of high probabilities; severe underestimates of low proportions	Fischhoff et al. 1977; Lichtenstein et al. 1982
Reluctance to report extreme events	
Weather forecasters provided accurate probabilities	Winkler and Murphy 1973
Poor estimates of expected severity	Dorris and Tabrizi 1977
Correlation of 0.72 between subjective and objective measures of injury frequency	Rethans 1980
Risk estimates lower for self than for others	Weinstein 1980, 1987
Risk estimates related to catastrophic potential, degree of control, familiarity	Lichtenstein et al. 1978
Evaluations of outcomes and probabilities are dependent	Weber 1994
Statistical Inference	
Conservative aggregation of evidence	Edwards 1966
Nearly optimal aggregation of evidence in naturalistic setting	Lehto et al. 2000
Failure to consider base rates	Tversky and Kahneman 1974
Base rates considered	Birnbaum and Mellers 1983 Koehler 1996
Overestimation of conjunctive events	Bar-Hillel 1973
Underestimation of disjunctive events	
Tendency to seek confirming evidence	Einhorn and Hogarth 1978; Baron 1985
Tendency to discount disconfirming evidence	
Tendency to ignore reliability of the evidence	Kahneman and Tversky 1973
Subjects considered variability of data when judging probabilities	Evans and Pollard 1985
People insensitive to information missing from fault trees	Fischhoff et al. 1978
Overconfidence in estimates	Fischhoff et al. 1977
Hindsight bias	Fishhoff 1982 Christensen-Szalanski and Willham 1991
Illusionary correlations	Tversky and Kahneman 1974
Gambler's fallacy	
Misestimation of covariance between items	Arkes 1981
Misinterpretation of regression to the mean	Tversky and Kahneman 1974

artifacts that do not reflect the true complexity of the world (Cohen 1993). From one such perspective, people deviate from Bayes' rule because it makes unrealistic assumptions about what is known or knowable. Simon (1955, 1983) makes a particularly compelling argument for the latter point of view. It also has been noted that researchers overreport findings of bias (Evans 1989; Cohen 1993).

There is an emerging body of literature that, on one hand, shows that deviations from Bayes' rule can in fact be justified in certain cases from a normative view and, on the other hand, shows that these deviations may disappear when people are provided richer information or problems in more natural contexts. For example, drivers performing a simulated passing task combined their own observations of the driving environment with imperfect information provided by a collision-warning system, as predicted by a distributed signal detection theoretic model of optimal team decision making (Lehto et al. 2000). Other researchers have pointed out that:

1. A tendency towards conservatism can be justified when evidence is not conditionally independence (Navon 1979).
2. Subjects do use base rate information and consider the reliability of evidence, in slightly modified experimental settings (Birnbaum and Mellers 1983; Koehler 1996).
3. A tendency to seek out confirming evidence can offer practical advantages (Cohen 1993) and may reflect cognitive failures, due to a lack of understanding of how to falsify hypotheses, rather than entirely a motivational basis (Klayman and Ha 1987; Evans 1989).
4. Subjects prefer stating subjective probabilities with vague verbal expressions rather than precise numerical values (Wallsten et al. 1993), demonstrating that they are not necessarily overconfident in their predictions.*
5. There is evidence that the hindsight bias can be moderated by familiarity with both the task and the type of outcome information provided (Christensen-Szalanski and Willham 1991).

Koehler (1996) provides a particularly compelling reexamination of the base rate fallacy. He concludes that the literature does not support the conventional wisdom that people routinely ignore base rates. To the contrary, he states that base rates are almost always used and that their degree of use depends on task structure and representation as well as their reliability compared to other sources of information. Because such conflicting conclusions can be obtained, depending upon the setting in which human decision making is observed, Koehler and researchers in the field of naturalistic decision making (Klein 1998; Klein et al. 1993) strongly emphasize the need to conduct ecologically valid research in rich realistic decision environments.

4.1.2. Heuristics and Biases

Tversky and Kahneman (1973, 1974) made a key contribution to the field when they showed that many of the above-mentioned discrepancies between human estimates of probability and Bayes' rule could be explained by the use of three heuristics.† The three heuristics they proposed were those of representativeness, availability, and anchoring and adjustment.

The *representativeness* heuristic holds that the probability of an item *A* belonging to some category *B* is judged by considering how representative *A* is of *B*. For example, a person is typically judged more likely to be a librarian than a farmer when described as "A meek and tidy soul, who has a desire for order and structure and a passion for detail." Application of this heuristic will often lead to good probability estimates but can lead to systematic biases. Tversky and Kahneman (1974) give several examples of such biases. In each case, representativeness influenced estimates more than other, more statistically oriented information. In the first study, subjects ignored base rate information (given by the experimenter) about how likely a person was to be either a lawyer or an engineer. Their judgments seemed to be based entirely on how representative the description seemed to be of either occupation. Tversky and Kahneman (1983) found people overestimated conjunctive probabilities in a similar experiment. Here, after being told that "Linda is 31 years old, single, outspoken, and very bright," most subjects said it was more likely she was both a bank teller and active as a feminist than simply a bank teller. In a third study, most subjects felt that the probability of more than 60% male births on a given day was about the same for both large and small hospitals (Tversky and Kahneman 1974). Apparently, the subjects felt large and small hospitals were equally representative of the population.

Other behaviors explained in terms of representativeness by Tversky and Kahneman included gambler's fallacy, insensitivity to predictability, illusions of validity, and misconceptions of statistical regression to the mean. With regard to gambler's fallacy, they note that people may feel long sequences of heads or tails when flipping coins are unrepresentative of normal behavior. After a sequence of heads, a tail therefore seems more representative. Insensitivity to predictability refers to a tendency for people to predict future performance without considering the reliability of the information they base the prediction upon. For example, a person might expect an investment to be profitable solely on the basis of a favorable description without considering whether the description has any predictive value. In other words, a good description is believed to be representative of high profits, even if it states nothing about profitability. The illusion of validity occurs when people use highly correlated evidence to make a conclusion. Despite the fact that the evidence is redundant, the

* It is interesting to note that Dawes and Mulford (1996) claim that the empirical support for the overconfidence effect is inadequate and logically flawed.

† It is important to point out that heuristic reasoning can lead to excellent results.

presence of many representative pieces of evidence increases confidence greatly. Misconception of regression to the mean occurs when people react to unusual events and then infer a causal linkage when the process returns to normality on its own. For example, a manager might incorrectly conclude that punishment works after seeing that unusually poor performance improves to normal levels following punishment. The same manager might also conclude that rewards don't work after seeing that unusually good performance drops after receiving a reward.

The availability heuristic holds that the probability of an event is determined by how easy it is to remember the event happening. Tversky and Kahneman state that perceived probabilities will, therefore, depend on familiarity, salience, effectiveness of memory search, and imaginability. The implication is that people will judge events as more likely when the events are familiar, highly salient (such as an airplane crash), or easily imaginable. Events also will be judged more likely if there is a simple way to search memory. For example, it is much easier to search for words in memory by the first letter rather than the third letter. It is easy to see how each of the above items impacting the availability of information can influence judgments. Biases should increase when people lack experience or when their experiences are too focused.

Anchoring and adjustment holds that people start from some initial estimate and then adjust it to reach some final value. The point initially chosen has a major impact on the final value selected when adjustments are insufficient. Tversky and Kahneman refer to this source of bias as an anchoring effect. They show how this effect can explain under- and overestimates of disjunctive and conjunctive events. This happens if the subject starts with a probability estimate of a single event. The probability of a single event is of course less than that for the disjunctive event and greater than that for the conjunctive event. If adjustment is too small, then under- and overestimates respectively occur for the disjunctive and conjunctive events. Tversky and Kahneman also discuss how anchoring and adjustment may cause biases in subjective probability distributions. Hogarth and Einhorn (1992) present an anchoring and adjustment model of how people update beliefs that explains a number of ordering effects, such as the primacy and recency effects. This latter model holds that the degree of belief in a hypothesis after collecting k pieces of evidence can be described as follows:

$$S_k = S_{k-1} + w_k[s(x_k) - R] \tag{26}$$

where S_k is the degree of belief after collecting k pieces of evidence, S_{k-1} is the anchor or prior belief, w_k is the adjustment weight for the kth piece of evidence, $s(x_k)$ is the subjective evaluation of the kth piece of evidence, and R is the reference point against which the kth piece of evidence is compared. In evaluation tasks, $R = 0$. This corresponds to the case where evidence is either for or against a hypothesis.* For estimation tasks, $R \neq 0$. The different values of R result in an additive model for evaluation tasks and an averaging model for estimation tasks. Also, if the quantity, $s(x_k) - R$, is evaluated for several pieces of evidence at a time, the model predicts primacy effects. If single pieces of evidence are individually evaluated in a step-by-step sequence, recency effects become more likely.

The notion of heuristics and biases has had a particularly formative influence on decision theory. A substantial recent body of work has emerged that focuses on applying research on heuristics and biases (Kahneman et al. 1982; Heath et al. 1994). Applications include medical judgment and decision making, affirmative action, education, personality assessment, legal decision making, mediation, and policy making. It seems clear that this approach is excellent for describing many general aspects of decision making in the real world. However, research on heuristics and biases has been criticized as being pretheoretical (Slovic et al. 1977). Koehler (1996) also points out that efforts to confirm the representativeness heuristic has contributed to overselling of the ''base rate'' bias. Other views of human judgment are expanded upon below.

4.1.3. Selective Processing of Information

Evans (1989) argues that factors which cause people to process information in a selective manner or attend to irrelevant information are the major cause of biases in human judgment. Factors assumed to influence selective processing include the availability, vividness, and relevance of information, and working memory limitations. The notion of availability refers to the information actually attended to by a person while performing a task. Evans's model assumes that relevant information elements are determined during a heuristic, preattentional stage. This stage is assumed to involve unconscious processes and is influenced by stimulus salience (or vividness) and the effects of prior knowledge.

In the next stage of his model, inferences are drawn from the selected information. This is done using rules for reasoning and action developed for particular types of problems. Working memory

*It is easy to see that Eq. (26) approximates the log-odds form of Bayes' rule where evidence for or against the hypothesis is additively combined.

influences performance at this stage by limiting the amount of information that can be consciously attended to while performing a task. The knowledge used during the inference process might be organized in schemas that are retrieved from memory and fit to specific problems (Cheng and Holyoak 1985). Support for this latter conclusion is provided by studies showing that people are able to develop skills in inference tasks but may fail to transfer these skills (inference related) from one setting to another. Evans also provides evidence that prior knowledge can cause biases when it is inconsistent with provided information and that improving knowledge can reduce or eliminate biases.

Evans's model of selective processing of information is consistent with other explanations of biases. Among such explanations, information overload has been cited as a reason for impaired decision making by consumers (Jacoby 1977). The tendency of highly salient stimuli to capture attention during inference tasks has also been noted by several researchers (Nisbett and Ross 1980; Payne 1980). Nisbett and Ross suggest that vividness of information is determined by its emotional content, concreteness and imagability, and temporal and spatial proximity. As noted by Evans, these factors have also been shown to affect the memorability of information. This provides a plausible explanation of both the availability heuristic and the experimental results mentioned earlier regarding biases in risk perceptions.

4.1.4. Models of Human Judgment

A number of approaches have been developed for mathematically describing human judgments. These approaches include social judgment theory, policy capturing, multiple-cue probability learning models, information integration theory, and conjoint measurement approaches.

Social judgment theory (SJT) implements an ecological approach for explaining how environmental cues are related to psychological responses (Hammond et al. 1975; Hammond 1993; Brehmer and Joyce 1988). The approach can be traced back to the Brunswick lens model (Brunwick 1952), which describes human judgments in terms of perceived environmental cues. Emphasis is placed on performing experiments where information cues reflect the statistical characteristics of the real world. Policy-capturing models are also derived from the lens model and have been applied to a wide number of real-world applications to describe expert judgments (Brehmer and Joyce 1988). For example, policy-capturing models have been applied to describe software selection by management information system managers (Martocchio et al. 1993), medical decisions (Brehmer and Joyce 1988), and highway safety (Hammond 1993). As mentioned earlier with regard to preference assessment, linear or nonlinear forms of regression are used in this approach to relate judgments to environment cues. These equations provide surprisingly good fits to expert judgments. In fact, there is evidence, and consequently much debate, over whether the models can actually do better than experts on many judgment tasks (Slovic et al. 1977; Brehmer 1981; Kleinmuntz 1984).

Cognitive continuum theory (Hammond 1980) builds upon Brunswick's earlier work by distinguishing judgments on a cognitive continuum varying from highly intuitive decisions to highly analytical decisions. Hammond (1993) summarizes earlier research showing that task characteristics cause decision makers to vary on this continuum. A tendency towards analysis increases, and reliance on intuition decreases, when (1) the number of cues increases, (2) cues are measured objectively instead of subjectively, (3) cues are of low redundancy, (4) decomposition of the task is high, (5) certainty is high, (6) cues are weighted unequally in the environmental model, (7) relations are nonlinear, (8) an organizing principle is available, (9) cues are displayed sequentially instead of simultaneously, and (10) the time period for evaluation is long. Intuitive methods can be better than analytical methods in some situations (Hammond et al. 1987).

Multiple cue probability learning models extend the lens model to the psychology of learning (Brehmer and Joyce 1988). Research on multiple-cue probability learning has provided valuable insight into factors affecting learning of inference tasks. One major finding is that providing cognitive feedback about cues and their relationship to the inferred effects leads to quicker learning than feedback about outcomes (Balzer et al. 1989). Stevenson et al. (1993) summarize a number of other findings, including that (1) subjects can learn to use valid cues, even when they are unreliable, (2) subjects are better able to learn linear relationships than nonlinear or inverse relationships, (3) subjects do not consider redundancy when using multiple cues, (4) source credibility and cue validity are considered, and (5) the relative effectiveness of cognitive and outcome feedback depends on the formal, substantive, and contextual characteristics of the task.

Information integration theory (Anderson 1981) takes a somewhat different approach than SJT or the lens model to describe how cue information is used when making judgments. A major deviation is that information-integration theory emphasizes the use of factorial experimental designs where cues are systematically manipulated. The goal of this approach is to determine first how people scale cues when determining their subjective values, and second how these scaled values are combined to form overall judgments. Various functional forms of how information is integrated are considered, including additive and averaging functions. A substantial body of research follows this approach to test various ways people might combine probabilistic information. A primary conclusion is that people tend to integrate information using simple averaging, adding, subtracting, and multiplying models. Conjoint measurement approaches (Wallsten 1972, 1976), in particular, provide a convenient way of both scaling subjective values assigned to cues and testing different functional forms describing how these values are combined to develop global judgments. By applying this approach, Wallsten (1976) was able to model primacy and recency effects.

4.1.5. Debiasing Human Judgments

The notion that many biases (or deviations from normative models) in statistical estimation and inference can be explained has led researchers to consider the possibility of debiasing human judgments (Keren 1990). Part of the issue is that heuristics often work very well. It seems logical that biases based on both the availability and representativeness heuristics might be reduced if people were provided more information. As discussed earlier in Section 4.1.1, there is evidence that biases can be moderated by familiarity with both the task and the type of outcome information provided. However, debiasing research has provided mixed results. Many biases, such as optimistic beliefs regarding health risks, have been difficult to modify (Weinstein and Klein 1995). People show a tendency to seek out information that supports their personal views (Weinstein 1979) and are quite resistant to information that contradicts strongly held beliefs (Nisbett and Ross 1980; McGuire 1966). Evans (1989) concludes that "pre-conceived notions are likely to prejudice the construction and evaluation of arguments."

Other evidence shows that experts may have difficulty providing accurate estimates of subjective probabilities even when they receive feedback. For example, many efforts to reduce both overconfidence in probability estimates and the hindsight bias have been unsuccessful (Fischhoff 1982). One problem is that people may not pay attention to feedback (Fischhoff and MacGregor 1982). They also may only attend to feedback that supports their hypothesis, leading to poorer performance and at the same time greater confidence (Einhorn and Hogarth 1978). Efforts to reduce confirmation biases through training have also in general been unsuccessful (Evans 1989).

On the positive side, there is evidence that providing feedback on the accuracy of weather forecasts may help weather forecasters (Winkler and Murphy 1973). There is also some evidence that people can learn to perform statistical reasoning more accurately after training in statistics (Fong et al. 1986). Failure to consider sample size was significantly reduced after training. Another study showed that asking people to write down reasons for and against their estimates of probabilities improved calibration and reduced overconfidence (Koriat et al. 1980). There is evidence that overconfidence is reduced when decision makers represent subjective probabilities verbally (Zimmer 1983; Wallsten et al. 1993). Conservatism, or the failure to modify probabilities adequately after obtaining evidence, was also reduced in Zimmer's study.

The conclusion is that debiasing human judgments is difficult but not impossible. Some perspective can be obtained by considering that most studies showing biases have focused on statistical inference and generally involved people not particularly knowledgeable about statistics, who are not using decision aids such as computers or calculators. It naturally may be expected that people will perform poorly on such tasks, given their lack of training and forced reliance on mental calculations (Winterfeldt and Edwards 1986). The finding that people can improve their abilities on such tasks after training in statistics is particularly telling, but also encouraging. Another encouraging finding is that biases are occasionally reduced when people process information verbally instead of numerically. This result might be expected, given that most people are more comfortable with words than numbers.

4.2. Preference and Choice

Much of the research on human preference and choice has focused on comparing observed preferences to the predictions of subjective utility theory (SEU) (Goldstein and Hogarth 1997). Early work, examining SEU as a descriptive theory, drew generally positive conclusions. However, it soon became apparent that people's preferences for risky or uncertain alternatives often violated basic axioms of SEU theory. The finding that people's preferences change when the outcomes are framed in terms of costs, as opposed to benefits, has been particularly influential. Several other common deviations from SEU have been observed. One potentially serious deviation is that preferences can be influenced by sunk costs or prior commitment to a particular alternative. Preferences change over time and may depend upon which alternatives are being compared, or even the order in which they are compared. The regret associated with making the "wrong" choice seems to play a major role when people compare alternatives. Accordingly, the satisfaction people derive from obtaining particular outcomes after making a decision is influenced by positive and negative expectations prior to making the decision. Other research on human preference and choice has shown that people choose between and apply different decision strategies, depending upon the cognitive effort required to apply a decision strategy successfully, the needed level of accuracy, and time pressure. Certain strategies are more likely than others to lead to choices consistent with those prescribed by SEU theory.

Alternative models, such as prospect theory and random utility theory, were consequently developed in order to explain human preferences under risk or uncertainty.* The following discussion will first summarize some common violations of the axioms underlying SEU theory before moving on to framing effects and preference reversals. Attention will then shift to models of choice and preference. The latter discussion will begin with prospect theory before addressing other models of labile or conditional preferences. Decision-making strategies, and how people choose between them, will be covered in Section 6.

4.2.1. *Violation of Rationality Axioms*

Several studies have shown that people's preferences between uncertain alternatives can be inconsistent with the axioms underlying subjective expected utility (SEU) theory. One fundamental violation of the assumptions is that preferences can be intransitive (Tversky 1969; Budescu and Weiss 1987). Also, as mentioned in the previous section, subjective probabilities may depend upon the values of consequences (violating the independence axiom) and, as discussed in the next section, the framing of a choice can impact preference. Another violation is given by the Myers effect (Myers et al. 1965), where preference reversals between high (H) and low (L) variance gambles can occur when the gambles are compared to a certain outcome, depending upon whether the certain outcome is positive (H preferred to L) or negative (L preferred to H). This latter effect violates the assumption of independence because the ordering of the two gambles depends on the certain outcome.

Another commonly cited violation of SEU theory is that people show a tendency towards uncertainty avoidance which can lead to behavior inconsistent with the "sure-thing" axiom. The Ellsburg and Allais paradoxes (Ellsburg 1961; Allais 1953) both involve violations of the sure-thing axiom (see Table 2) and seem to be caused by people's desire to avoid uncertainty. The Allais paradox is illustrated by the following set of gambles. In the first gamble, a person is asked to choose between gambles $A1$ and $B1$, where:

Gamble $A1$ results in \$1 million for sure. Gamble $B1$ results in \$2.5 million with a probability of 0.1, \$1 million with a probability of 0.89, and \$0 with a probability of 0.01.

In the second gamble, the person is asked to choose between gambles *A2 and* $B2$, where: $A2$ results in \$1 million with a probability of 0.11 and \$0 with a probability of 0.89. Gamble $B2$ results in \$2.5 million with a probability of 0.1 and \$0 with a probability of 0.9.

Most people prefer gamble $A1$ to $B1$ and gamble $B2$ to $A2$. It is easy to see that this set of preferences violates expected utility theory. First, if $A1 > B1$, then $u(A1) > u(B1)$, meaning that: $u(\$1 \text{ million}) > 0.1u(\$2.5 \text{ million}) + 0.89u(\$1 \text{ million}) + 0.01u(\$0)$. If a utility of 0 is assigned to receiving \$0 and a utility of 1 to receiving \$2.5 million, then $u(\$1 \text{ million}) > 1/11$. However, from the preference $A2 > B2$, it follows that $u(\$1 \text{ million}) < 1/11$. Obviously, no utility function can satisfy this requirement of assigning a value both greater than and less than 1/11 to \$1 million.

As noted by Savage (1954), the above set of gambles can be reframed in a way that shows that these preferences violate the sure-thing principle. After doing so, Savage found that his initial tendency towards choosing $A1$ over $B1$ and $A2$ over $B2$ disappeared. As noted by Stevenson et al. (1993), this example is one of the first cited cases of a preference reversal caused by reframing a decision, the topic discussed below.

4.2.2. *Framing of Decisions and Preference Reversals*

A substantial body of research has shown that people's preferences can shift dramatically depending upon the way a decision is represented. The best-known work on this topic was conducted by Tversky and Kahneman (1981), who showed that preferences between medical intervention strategies changed dramatically depending upon whether the outcomes were posed as losses or gains. The following question, worded in terms of benefits, was presented to one set of subjects:

> Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:
>
> If Program A is adopted, 200 people will be saved.
>
> If Program B is adopted, there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved.
>
> Which of the two programs would you favor?

* Yates (1992) and Singleton and Hovden (1987) are useful sources for the reader interested in additional details on risk perception, risk acceptability, and risk-taking behavior. Section 5.1.1 is also relevant to this topic.

The results showed that 72% of subjects preferred program A. The second set of subjects was given the same cover story, but worded in terms of costs, as given below:

If Program C is adopted, 400 people will die.

If Program D is adopted, there is a 1/3 probability that nobody will die, and a 2/3 probability that 600 people will die.

Which of the two programs would you favor?

The results now showed that 78% of subjects preferred program D. Since program D is equivalent to B and Program A is equivalent to C, the preferences for the two groups of subjects were strongly reversed. Tversky and Kahneman concluded that this reversal illustrated a common pattern in which choices involving gains are risk averse and choices involving losses are risk seeking. The interesting result was that the way the outcomes were worded caused a shift in preference for identical alternatives. Tversky and Kahneman called this tendency the *reflection effect*. A body of literature has since developed showing that the framing of decisions can have practical effects for both individual decision makers (Kahneman et al. 1982; Heath et al. 1994) and group decisions (Paese et al. 1993). On the other hand, recent research shows that the reflection effects can be reversed by certain outcome wordings (Kuhberger 1995); more importantly, Kuhberger provides evidence that the reflection effect observed in the classic experiments can be eliminated by fully describing the outcomes (i.e., referring to the above paragraph, a more complete description would state, "If Program C is adopted, 400 people will die AND 200 WILL LIVE").

Other recent research has explored the theory that perceived risk and perceived attractiveness of risky outcomes are psychologically distinct constructs (Weber et al. 1992). In the latter study, it was concluded that perceived risk and attractiveness are "closely related, but distinct phenomena." Related research has shown weak negative correlations between the perceived risk and value of indulging in alcohol-related behavior for adolescent subjects (Lehto et al. 1994). This latter study also showed that the rated propensity to indulge in alcohol-related behavior was strongly correlated with perceived value ($R = 0.8$), but weakly correlated with perceived risk ($R = -0.15$). Both findings are consistent with the theory that perceived risk and attractiveness are distinct constructs, but the latter finding indicates that perceived attractiveness may be the better predictor of behavior. Lehto et al. conclude that intervention methods attempting to lower preferences for alcohol-related behavior should focus on lowering perceived value rather than on increasing perceived risk.

4.2.3. Prospect Theory

Prospect theory (Kahneman and Tversky 1979) attempts to account for behavior not consistent with the SEU model by including the framing of decisions as a step in the judgment of preference between risky alternatives. Prospect theory assumes that decision makers tend to be risk averse with regard to gains and risk seeking with regard to losses. This leads to a value function that disproportionately weights losses. As such, the model is still equivalent to SEU, assuming a utility function expressing mixed risk aversion and risk seeking. Prospect theory, however, assumes that the decision maker's reference point can change. With shifts in the reference point, the same returns can be viewed as either gains or losses.* This latter feature of prospect theory, of course, is an attempt to account for the framing effect discussed above. Prospect theory also deviates significantly from SEU theory in the way probabilities are addressed. To describe human preferences more closely, perceived values are weighted by a function $\pi(p)$, instead of the true probability, p. Compared to the untransformed form of p, $\pi(p)$ overweights very low probabilities and underweights moderate and high probabilities. The function $\pi(p)$ is also generally assumed to be discontinuous and poorly defined for probability values close to 0 or 1.

Prospect theory assumes that the choice process involves an editing phase and an evaluation phase. The editing phase involves reformulation of the options to simplify subsequent evaluation and choice. Much of this editing process is concerned with determining an appropriate reference point in a step called coding. Other steps that may occur include the segregation of riskless components of the decision, combining probabilities for events with identical outcomes, simplification by rounding off probabilities and outcome measures, and search for dominance. In the evaluation phase, the perceived values are then weighed by the function $\pi(p)$. The alternative with the greatest weighed value is then selected. Several other modeling approaches that differentially weigh utilities in risky decision making

*The notion of a reference point against which outcomes are compared has similarities to the notion of making decisions on the basis of regret (Bell 1982). Regret, however, assumes comparsion to the best outcome. The notion of different reference points also is related to the well-known trend that buying and selling price of assets often differ for a decision maker (Raiffa 1968).

have been proposed (Goldstein and Hogarth 1998). As in prospect theory, such models often assume that the subjective probabilities, or decision weights, are a function of outcome sign (i.e. positive, neutral, or negative), rank (i.e., 1st, 2nd, etc), or magnitude. Other models focus on display effects (i.e., single-stage vs. multistage arrangements) and distribution effects (i.e., two outcome lotteries vs multiple-outcome lotteries). Prospect theory and other approaches also address how the value or utility of particular outcomes can change between decision contexts, as discussed below.

4.2.4. Labile Preferences

There is no doubt that human preferences often change after receiving some outcome. After losing money, an investor may become risk averse. In other cases, an investor may escalate her commitment to an alternative after an initial loss, even if better alternatives are available. From the most general perspective, any biological organism becomes satiated after satisfying a basic need, such as hunger. Preferences also change over time or between decision contexts. For example, a 30-year-old decision maker considering whether to put money into a retirement fund may currently have a very different utility function than at retirement. The latter case is consistent with SEU theory but obviously complicates analysis.

Economists and behavioral researchers have both focused on mathematically modeling choice processes to explain intransitive or inconsistent preference orderings of alternatives (Goldstein and Hogarth 1997). Game theory provides interesting insight into this issue. From this perspective, preferences of the human decision maker are modeled as the collective decisions obtained by a group of internal agents, or selves, each of which is assumed to have distinct preferences (see Elster 1986). Intransitive preferences and other violations of rationality on the part of the human decision maker then arise from interactions between competing selves.* Along these lines, Ainslie (1975) proposed that impulsive preference switches (often resulting in risky or unhealthy choices) arise as the outcome of a struggle between selves representing conflicting short-term and long-term interests, respectively.

Another area of active research has focused on how experiencing outcomes can cause shifts in preference. One robust finding is that people tend to be more satisfied if an outcome exceeds their expectations and less satisfied if it does not (i.e., Feather 1966; Connolly et al. 1997). Expectations therefore provide a reference point against which obtained outcomes are compared. Numerous studies have also shown that people in a wide variety of settings often consider sunk costs when deciding whether to escalate their commitment to an alternative by investing additional resources (Arkes and Blumer 1985). From the perspective of prospect theory, sunk costs cause people to frame their choice in terms of losses instead of gains, resulting in risk-taking behavior and consequently escalating commitment. Other plausible explanations for escalating commitment include a desire to avoid waste or to avoid blame for an initially bad decision to invest in the first place. Interestingly, some recent evidence suggests that people may deescalate commitment in response to sunk costs (Heath 1995). The latter effect is also contrary to classical economic theory, which holds that decisions should be based solely on marginal costs and benefits. Heath explains such effects in terms of mental accounting. Escalation is held to occur when a mental budget is not set or expenses are difficult to track. Deescalation is held to occur when people exceed their mental budget, even if the marginal benefits exceed the marginal costs.

Other approaches include value or utility as random variables within models of choice to explain intransitive or inconsistent preference orderings of alternatives. Random utility models (Iverson and Luce 1998) describe the probability $P_{a,A}$ of choosing a given alternative a from a set of options A as

$$P_{a,A} = \text{Prob}(U_a \geq U_b, \text{ for all } b \text{ in } A) \tag{27}$$

where U_a is the uncertain utility of alternative a and U_b is the uncertain utility of alternative b. The most basic random utility models assign a utility to each alternative by sampling a single value from some known distribution. The sampled utility of each alternative then remains constant throughout the choice process. Basic random utility models can predict a variety of preference reversals and intransitive preferences for single and multiple attribute comparisons of alternatives (i.e., Tverski 1972).

Sequential sampling models extend this approach by assuming preferences can be based on more than one observation. Preferences for particular alternatives are accumulated over time, by integrating

*As discussed further in Section 6, group decisions, even though they are made by rational members, are subject to numerous violations of rationality. For example, consider the case where the decision maker has three selves that are, respectively, risk averse, risk neutral, and risk seeking. Assume that the decision maker is choosing between alternatives A, B, and C. Suppose the risk-averse self rates the alternatives in the order A, B, C; the risk-neutral self rates them in the order B, C, A; and the risk-seeking self rates them in the order C, A, B. Also, assume the selves are equally powerful. Then two of the three agents always agree that $A > B$, $B > C$, and $C > A$. This ordering is, of course, nontransitive.

or otherwise summing the sampled utilities. The utility of an alternative, at a particular time, is proportional to the latter sum. A choice is made when the summed preferences for a particular alternative exceed some threshold, which itself may vary over time or depend on situational factors (Wallsten 1995; Busemeyer and Townsend 1993). It is interesting to observe that sequential sampling models can explain speed accuracy trade-offs in signal-detection tasks (Stone 1960), as well as shifts in preferences due to time pressure (Wallsten 1995; Busemeyer and Townsend 1993), if it is assumed that people adjust their threshold downwards under time pressure. That is, under time pressure, people sample less information before making a choice. The following section will further explore how and why decision strategies might change over time and between decision contexts.

5. DYNAMIC AND NATURALISTIC DECISION MAKING

In dynamic decision making, actions taken by a decision maker are made sequentially in time. Taking actions can change the environment, resulting in a new set of decisions. The decisions might be made under time pressure and stress, by groups or by single decision makers. This process might be performed on a routine basis or might involve severe conflict. For example, either a group of soldiers or an individual officer might routinely identify marked vehicles as friends or foes. When a vehicle has unknown or ambiguous marking, the decision changes to a conflict driven process. Naturalistic decision theory has emerged as a new field that focuses on such decisions in real-world environments (Klein et al. 1993; Klein 1998). The notion that most decisions are made in a routine, nonanalytical way is the driving force of this approach.* Areas where such behavior seems prominent include juror decision making, troubleshooting of complex systems, medical diagnosis, management decisions, and numerous other examples.

The following discussion will first address models of dynamic and naturalistic decision making. These models both illustrate naturalistic decision-making strategies and explain their relation to experience and task familiarity. A brief discussion will also be provided on teams and team leadership, in naturalistic settings. Attention will then shift to the issue of time pressure and stress and how this factor influences performance in naturalistic decision making.

5.1. Naturalistic Decision Making

In recent years, it has been recognized that decision making in natural environments often differs greatly between decision contexts (Beach 1993; Hammmond 1993). In addressing this topic, the involved researchers often question the relevance and validity of both classical decision theory and behavioral research not conducted in real-world settings (Cohen 1993). Numerous naturalistic models have been proposed (Klein et al. 1993). These models assume that people rarely weigh alternatives and compare them in terms of expected value or utility. Each model is also descriptive rather than prescriptive. Perhaps the most general conclusion that can be drawn from this work is that people use different decision strategies, depending upon their experience, the task and the decision context. Several of the models also postulate that people choose between decision strategies by trading off effectiveness against the effort required.

The following discussion will briefly review seven modeling perspectives that fit into this framework: (1) levels of task performance (Rasmussen 1983), (2) recognition-primed decisions (Klein 1989), (3) image theory (Beach 1990), (4) contingent decision making (Payne et al. 1993), (5) dominance structuring (Montgomery 1989), (6) explanation-based decision making (Pennington and Hastie 1988), and (7) shared mental models and awareness. Attention will then shift to leadership and its impact on team performance in naturalistic settings.

5.1.1. Levels of Task Performance

There is growing recognition that most decisions are made on a routine basis in which people simply follow past behavior patterns (Rasmussen 1983; Beach 1993; Svenson 1990). Rasmussen (1983) follows this approach to distinguish among skill-based, rule-based, and knowledge-based levels of task performance. Lehto (1991) further considers judgment-based behavior as a fourth level of performance.

Performance is said to be at either a skill-based or a rule-based level when tasks are routine in nature. Skill-based performance involves the smooth, automatic flow of actions without conscious decision points. As such, skill-based performance describes the decisions made by highly trained operators performing familiar tasks. Rule-based performance involves the conscious perception of environmental cues, which trigger the application of rules learned on the basis of experience. As such, rule-based performance corresponds closely to recognition-primed decisions (Klein 1989). The

*Drucker (1985), in discussing ways of improving the effectiveness of executive decision makers, emphasizes the importance of establishing a generic principle or policy that can be applied to specific cases in a routine way. This recommendation is interesting because it prescribes a naturalistic form of behavior.

knowledge-based level of performance is said to occur during learning or problem-solving activity during which people cognitively simulate the influence of various actions and develop plans for what to do. The judgment-based level of performance occurs when affective reactions of a decision maker cause a change in goals or priorities between goals (Janis and Mann 1977; Etzioni 1988; Lehto 1991). Distinctive types of errors in decision making occur at each of the four levels (Reason 1989; Lehto 1991).

At the skill-based level, errors occur due to perceptual variability and when people fail to shift up to rule-based or higher levels of performance. At the rule-based level, errors occur when people apply faulty rules or fail to shift up to a knowledge-based level in unusual situations where the rules they normally use are no longer appropriate. The use of faulty rules leads to an important distinction between running and taking risks. Along these lines, Wagenaar (1992) discusses several case studies in which people following risky forms of behavior do not seem to be consciously evaluating the risk. Drivers, in particular, seem to habitually take risks. Wagenaar explains such behavior in terms of faulty rules derived on the basis of benign experience. In other words, drivers get away with providing small safety margins most of the time and consequently learn to run risks on a routine basis. Drucker (1985) points out several cases where organizational decision makers have failed to recognize that the generic principles they used to apply were no longer appropriate, resulting in catastrophic consequences.

At the knowledge-based level, errors occur because of cognitive limitations or faulty mental models or when the testing of hypotheses cause unforeseen changes to systems. At the judgment-based levels, errors (or violations) occur because of inappropriate affective reactions, such as anger or fear (Lehto 1991). As noted by Isen (1993), there also is growing recognition that positive affect can influence decision making. For example, positive affect can promote the efficiency and thoroughness of decision making, but may cause people to avoid negative materials. Positive affect also seems to encourage risk-averse preferences. Decision making itself can be anxiety provoking, resulting in violations of rationality (Janis and Mann 1977).

A recent study involving drivers arrested for drinking and driving (McKnight et al. 1995) provides an interesting perspective on how the sequential nature of naturalistic decisions can lead people into traps. The study also shows how errors can occur at multiple levels of performance. In this example, decisions made well in advance of the final decision to drive while impaired played a major role in creating situations where drivers were almost certain to drive impaired. For instance, the driver may have chosen to bring along friends and therefore have felt pressured to drive home because the friends were dependent upon him or her. This initial failure by drivers to predict the future situation could be described as a failure to shift up from a rule-based level to a knowledge-based level of performance. In other words, the driver never stopped to think about what might happen if he or she drank too much. The final decision to drive, however, would correspond to an error (or violation) at the judgment-based level if the driver's choice was influenced by an affective reaction (perceived pressure) to the presence of friends wanting a ride.

5.1.2. Recognition-Primed Decision Making

Klein (1989, 1998) developed the theory of recognition-primed decision making on the basis of observations of firefighters and other professionals in their naturalistic environments. He found that up to 80% of the decisions made by firefighters involved some sort of situation recognition, where the decision makers simply followed a past behavior pattern once they recognized the situation.

The model he developed distinguishes between three basic conditions. In the simplest case, the decision maker recognizes the situation and takes the obvious action. A second case occurs when the decision maker consciously simulates the action to check whether it should work before taking it. In the third and most complex case, the action is found to be deficient during the mental simulation and is consequently rejected. An important point of the model is that decision makers don't begin by comparing all the options. Instead, they begin with options that seem feasible based upon their experience. This tendency, of course, differs from the SEU approach but is comparable to applying the satisficing decision rule (Simon 1955) discussed earlier.

Situation assessment is well recognized as an important element of decision making in naturalistic environments (Klein et al. 1993). Recent research by Klein and his colleagues has examined the possibility of enhancing situation awareness through training (Klein and Wolf 1995). Klein and his colleagues have also applied methods of cognitive task analysis to naturalistic decision-making problems. In these efforts, they have focused on identifying (1) critical decisions, (2) the elements of situation awareness, (3) critical cues indicating changes in situations, and (4) alternative courses of action (Klein 1995). Accordingly, practitioners of naturalistic decision making tend to focus on process-tracing methods and behavioral protocols (Ericsson and Simon 1984) to document the processes people follow when they make decisions.*

* Goldstein and Hogarth (1998) describe a similar trend in judgment and decision-making research.

5.1.3. Image Theory

Image theory (Beach 1990) is a descriptive theory of decision making. Beach theorizes that knowledge used to make decisions falls into three categories: value images, trajectory images, and strategic images. The value image describes the decision maker's values, and principles; the trajectory image describes goals; the strategic image describes plans to attain the goals. He also theorizes that there are two types of decisions: adoption decisions and progress decisions. Adoption decisions first involve a screening process where alternatives are eliminated from consideration. The most promising alternative is then selected from the screened set. Progress decisions involve a comparison between goals and the expected result of choosing the alternative.

Two means of evaluating decisions are applied. One test compares the compatibility of the generated alternatives to value images, trajectory images, and strategic images. The profitability test is used to evaluate screened options further in adoption decisions when more than one option survives the initial screening. Beach (1993) argues strongly for the primacy of screening as a characteristic of most real-world decision-making activity.

5.1.4. Contingent Decision Making

The theory of contingent decision making (Beach and Mitchell 1978; Payne et al. 1993) is similar to image theory and cognitive continuum theory (see Section 4.1.4) in that it holds that people use different decision strategies, depending upon the characteristics of the task and the decision context. Payne et al. limit their modeling approach to tasks that require choices to be made (simple memory tasks are excluded from consideration). They also add the assumption that people make choices about how to make choices.*

Choices between decision strategies are assumed to be made rationally by comparing their cost (in terms of cognitive effort) against their benefits (in terms of accuracy). Cognitive effort and accuracy (of a decision strategy) are both assumed to depend upon task characteristics, such as task complexity, response mode, and method of information display. Cognitive effort and accuracy also are assumed to depend upon contextual characteristics, such as the similarity of the compared alternatives, attribute ranges and correlations, the quality of the considered options, reference points, and decision frames. Payne et al. place much emphasis on measuring the cognitive effort of different decision strategies in terms of the number of elemental information elements that must be processed for different tasks and contexts. They relate the accuracy of different decision strategies to task characteristics and contexts and also present research showing that people will shift decision strategies to reduce cognitive effort, increase accuracy, or in response to time pressure.

5.1.5. Dominance Structuring

Dominance structuring (Montgomery 1989) holds that decision making in real contexts involves a sequence of four steps. The process begins with a preediting stage in which alternatives are screened from further analysis. The next step involves selecting a promising alternative from the set of alternatives that survive the initial screening. A test is then made to check whether the promising alternative dominates the other surviving alternations. If dominance is not found, then the information regarding the alternatives is restructured in an attempt to force dominance. This process involves both the bolstering and deemphasizing of information in a way that eliminates disadvantages of the promising alternative.

5.1.6. Explanation-Based Decision Making

Explanation-based decision making (Pennington and Hastie 1986, 1988) assumes that people begin their decision-making process by constructing a mental model that explains the facts they have received. While constructing this explanatory model, people are also assumed to be generating potential alternatives to choose between. The alternatives are then compared to the explanatory model, rather than to the facts from which it was constructed.

Pennington and Hastie have applied this model to juror decision making and obtained experimental evidence that many of its assumptions seem to hold. They note that juror decision making requires consideration of a massive amount of data that is often presented in haphazard order over a long time period. Jurors seem to organize this information in terms of stories describing causation and intent. As part of this process, jurors are assumed to evaluate stories in terms of their uniqueness, plausibility, completeness, or consistency. To determine a verdict, jurors then judge the fit between choices provided by the trial judge and the various stories they use to organize the information.

*As such, the theory of contingent decision making directly addresses a potential source of conflict shown in the integrative model of decision making presented earlier (Figure 1). That is, it states that decision makers must choose between decision strategies when they are uncertain how to compare alternatives."

Jurors' certainty about their verdict is assumed to be influenced by both evaluation of stories and the perceived goodness of fit between the stories and the verdict.

5.1.7. *Shared Mental Models and Awareness*

Orasanu and Salas (1993) discuss two closely related frameworks for describing the knowledge used by teams in naturalistic settings. These are referred to as *shared mental models* and the *team mind*. The common element of these two frameworks is that the members of teams hold knowledge in common and organize it in the same way. Orasanu and Salas claim that this improves and minimizes the need for communication between team members, enables team members to carry out their functions in a coordinated way, and minimizes negotiation over who should do what at what time. Under emergency conditions, Orasanu and Salas claim there is a critical need for members to develop a shared situation model. As evidence for the notion of shared mental models and the team mind, the authors cite research in which firefighting teams and individual firefighters developed the same solution strategies for situations typical of their jobs.

This notion of shared mental models and the team mind can be related to the notion discussed earlier of schemas containing problem-specific rules and facts (Cheng and Holyoak 1985). It also might be reasonable to consider other team members as a form of external memory (Newell and Simon 1972). This approach would have similarities to Wegner's (1987) concept of transactive memory where people in a group know who has specialized information of one kind or another. Klein (1998) provides an interesting discussion of how this metaphor of the team mind corresponds to thinking by individuals. Teams, like people, have a working memory that contains information for a limited time, a long-term or permanent memory, and limited attention. Like people, they also filter out and process information and learn in many ways.

5.1.8. *Team Leadership*

Torrance (1953) describes retrospective accounts of military survivors lost behind enemy lines indicating that survival depended upon the leader's leadership skills. Important elements of leadership skills included keeping the members of the group focused on a common goal, making sure they knew what needed to be done, and keeping them informed of the current status. Related conclusions concerning the value of keeping people informed have been obtained in retrospective accounts of survivors of mining accidents (Mallet et al. 1993). Orasanu and Salas (1993) cite research in which captains of high-performing air crews explicitly stated more plans, strategies, and intentions to the other members of the crew. They also gave more warnings and predictions to the crew members. Orasanu and Salas cite other work showing that crews performed better with captains who were task oriented and had good personal skills. Performance dropped when captains had negative expressive styles and low task orientation.

A complementary literature has been developed on leadership theory (Chemers and Ayman 1993). Most of this research is based on leaders in organizational contexts. A sampling of factors which have been shown to be related to the effectiveness of leadership include legitimacy, charisma, individualized attention to group members, and clear definitions of goals. These results seem quite compatible with the above findings for leadership in naturalistic, dynamic contexts.

5.2. Time Pressure and Stress

Time pressure and stress are a defining characteristic of naturalistic decision making. Jobs requiring high levels of skill or expertise, such as firefighting, nursing, emergency care, and flying an airplane, are especially likely to involve high stakes, extreme time pressure, uncertainty, or risk to life. The effect of stressors, such as those mentioned above, on performance has traditionally been defined in terms of physiological arousal.* The Yerkes–Dodson law (Yerkes and Dodson 1908) states that the relation between performance and arousal is an inverted U. Either too much or too little arousal causes performance to drop. Too little arousal makes it difficult for people to maintain focused attention. Too much arousal results in errors, more focused attention (and filtering of low-priority information), reduced working memory capacity, and shifts in decision strategies.† One explanation of why performance drops when arousal levels are too high is that arousal consumes cognitive resources that could be allocated to task performance (Mandler 1979).

Time pressure is a commonly studied stressor assumed to impact decision making. Maule and Hockey (1993) note that people tend to filter out low-priority types of information, omit processing

*The general adaptation syndrome (Selye 1936, 1979) describes three stages of the human response to stressors. In simplified form, this sequence corresponds to (1) arousal, (2) resistance, and (3) exhaustion.

†The literature on stress and its effects on decision making will not be surveyed here. Books edited by Hamilton and Warburton (1979), Svenson and Maule (1993), Driskell and Salas (1996), and Flin et al. (1997) provide a good introduction to the area.

information, and accelerate mental activity when they are under time pressure. Variable state activation theory (VSAT) provides a potential explanation of the above effects in terms of a control model of stress regulation (Maule and Hockey 1993). Sequential sampling models provide a compatible perspective on how time pressure can cause changes in performance, such as speed–accuracy trade-offs (see Section 4.2.4). The two approaches are compatible, because VSAT provides a means of modeling how the decision thresholds used within a sequential sampling model might change as a function of time pressure. VSAT also proposes that disequilibriums between control processes and the demands of particular situations can lead to strong affective reactions or feelings of time pressure. Such reactions could, of course, lead to attentional narrowing or reduced working memory capacity and therefore result in poorer task performance. Alternatively, performance might change when decision thresholds are adjusted.

Time pressure also can cause shifts between the cognitive strategies used in judgment and decision-making situations (Payne et al. 1993; Maule and Hockey 1993; Edland and Svenson 1993). People show a strong tendency to shift to noncompensatory decision rules when they are under time pressure. This finding is consistent with contingency theories of strategy selection (Section 5.1.4). In other words, this shift may be justified when little time is available, because a noncompensatory rule can be applied more quickly. Compensatory decision rules also require more analysis and cognitive effort. Intuitive decision strategies require much less effort because people can rely on their experience or knowledge, and can lead to better decisions in some situations (Hammond et al. 1987). As Klein (1998) points out, stress should impact performance if people use analytical choice procedures.

Novices and experts in novel, unexpected, situations will lack domain experience and knowledge and therefore will have to rely on analytical choice procedures. Consequently, it is not surprising that time pressure and stress have a major negative impact on novice decision makers performing unfamiliar tasks. Interestingly, there is little evidence that stress or time pressure causes experienced personnel to make decision errors in real-world tasks (Klein 1996; Orasanu 1997). The latter finding is consistent with research indicating that experts rely on their experience and intuition when they are under stress and time pressure (Klein 1998). The obvious implication is that training and experience are essential if people are to make good decisions under time pressure and stress.

6. GROUP DECISION MAKING

Much research has been done over the past 25 years or so on decision making by groups and teams. Most of this work has focused on groups, as opposed to teams. In a team, it is assumed that the members are working toward a common goal and have some degree of inderdependence, defined roles and responsibilities, and task-specific knowledge (Orasanu and Salas 1993). Team performance is a major area of interest in the field of naturalistic decision theory (Klein et al. 1993; Klein 1998), as discussed earlier. Group performance has traditionally been an area of study in the fields of organizational behavior and industrial psychology. Traditional decision theory has also devoted some attention to group decision making (Raiffa 1968; Keeney and Raiffa 1976). The following discussion will first briefly discuss some of the ways that group decisions differ from those made by isolated decision makers who need to consider only their own preferences. That is, ethics and social norms play a much more prominent role when decisions are made by or within groups. Attention will then shift to group processes and how they affect group decisions. The last section will address methods of supporting or improving group decision making.

6.1. Ethics and Social Norms

When decisions are made by or within groups, a number of issues arise that have not been touched upon in the earlier portions of this chapter. To start, there is the complication that preferences may vary between members of a group. It often is impossible to maximize the preferences of all members of the group, meaning that trade-offs must be made and issues such as fairness must be addressed to obtain acceptable group decisions. Another complication is that the return to individual decision makers can depend on the actions of others. Game theory* distinguishes two common variations of this situation. In competitive games, individuals are likely to take "self-centered" actions that maximize their own return but reduce returns to other members of the group. Behavior of group members in this situation may be well described by the minimax decision rule discussed in Section 2.1.5. In cooperative games, the members of the group take actions that maximize returns to the group as a whole.

Members of groups may choose cooperative solutions that are better for the group as a whole for many different reasons (Dawes et al. 1988). Groups may apply numerous forms of coercion to punish members who deviate from the cooperative solutions. Group members may apply decision strategies such as reciprocal altruism. They also might conform because of their social conscience, a need for

* Friedman (1990) provides an excellent introduction to game theory.

self esteem, or feelings of group identity. Fairness considerations can in some case explain preferences and choices that seem to be in conflict with economic self-interest (Bazerman 1998). Changes in the status quo, such as increasing the price of bottled water immediately after a hurricane, may be viewed as unfair even if they are economically justifiable based on supply and demand. People are often willing to incur substantial costs to punish "unfair" opponents and reward their friends or allies. The notion that costs and benefits should be shared equally is one fairness-related heuristic people use (Messick 1991). Consistent results were found by Guth et al. (1982) in a simple bargaining game where player 1 proposes a split of a fixed amount of cash and player 2 either accepts the offer or rejects it. If player 2 rejects the offer, both players receive nothing. Classical economics predicts that player 2 will accept any positive amount (that is, player 2 should always prefer something to nothing). Consequently, player 1 should offer player 2 a very small amount greater than zero. The results showed that, contrary to predictions of classical economics, subjects tended to offer a substantial proportion of the cash (the average offer was 30%). Some of the subjects rejected positive offers. Others accepted offers of zero. Further research, as summarized by Bolton and Chatterjee (1996), confirms these findings that people seem to care about whether they receive their fair share.

Ethics clearly plays an important role in decision making. Some choices are viewed by nearly everyone as being immoral or wrong (i.e., violations of the law, dishonesty, and numerous other behaviors that conflict with basic societal values or behavioral norms). Many corporations and other institutions formally specify codes of ethics prescribing values such as honesty, fairness, compliance with the law, reliability, considerance or sensitivity to cultural differences, courtesy, loyalty, respect for the environment, and avoiding waste. It is easy to visualize scenarios, where it is in the best interest of a decision maker to choose economically undesirable options (at least in the short term) to comply with ethical codes. According to Kidder (1995), the "really tough choices . . . don't center on right versus wrong. They involve right versus right." Kidder refers to four dilemmas of right vs. right he feels qualify as paradigms: (1) truth vs. loyalty (i.e., whether to divulge information provided in confidence), (2) individual vs. community, (3) short term vs. long term, and (4) justice vs. mercy. At least three principles, which in some cases provide conflicting solutions, have been proposed for resolving ethical dilemmas. These include (1) utilitarianism, or selecting the option with the best overall consequences, (2) rule-based, or following a rule regardless of its current consequences (i.e., waiting for a stop light to turn green, even if no cars are coming), and (3) fairness, or doing what you would want others to do for you.

Numerous social dilemmas also occur in which the payoffs to each participant result in individual decision strategies harmful to the group as a whole. The tragedy of the commons (Hardin 1968) is illustrative of social dilemmas in general. For a recent example, discussed in detail by Baron (1998), consider the recent crash of the East Coast commercial fishing industry, brought about by overfishing. Here, the fishing industry as a whole is damaged by overfishing, but individual fishers gain a short-term advantage by catching as many fish as possible. Individual fishers may reason that if they don't catch the fish, someone else will. Each fisher attempts to catch as many fish as possible, even if this will cause the fish stocks to crash. Despite the fact that cooperative solutions, such as regulating the catch, are obviously better than the current situation, individual fishers continue to resist such solutions. Regulations are claimed to infringe on personal autonomy, to be unfair, or to based on inadequate knowledge.

Other similar examples include littering, wasteful use of natural resources, pollution, or social free riding. These behaviors can all be explained, in terms of the choices faced by the offending individual decision maker (Schelling 1978). Simply put, the individual decision maker enjoys the benefits of the offensive behavior, as small as they may be, but the costs are incurred by the entire group.

6.2. Group Processes

A large amount of research has focused on groups and their behavior. Accordingly, many models have been developed that describe how groups make decisions. A common observation is that groups tend to move through several phases as they go through the decision-making process (Ellis and Fisher 1994). One of the more classic models (Tuckman 1965) describes this process with four words: forming, storming, norming, and performing. Forming corresponds to initial orientation, storming to conflict, norming to developing group cohesion and expressing opinions, and performing to obtaining solutions. As implied by Tuckman's choice of terms, there is a continual interplay between socio-emotive factors and rational, task-oriented behavior throughout the group decision-making process. Conflict, despite its negative connotations, is a normal, expected aspect of the group decision process and can in fact serve a positive role (Ellis and Fisher 1994). The following discussion will first address causes and effects of group conflict. Attention will then shift to conflict resolution.

6.2.1. Conflict

Whenever people or groups have different preferences, conflict can occur. As pointed out by Zander (1994), conflict between groups becomes more likely when groups have fuzzy or potentially antagonistic roles, or when one group is disadvantaged (or perceives it is not being treated fairly). A lack of conflict-settling procedures and separation or lack of contact between groups can also contribute to conflict. Conflict becomes especially likely during a crisis and often escalates when the issues are perceived to be important, or after resistance or retaliation occurs. Polarization, loyalty to one's own group, lack of trust, and cultural and socioeconomic factors are often contributing factors to conflict and conflict escalation.

Ellis and Fisher (1994) distinguish between affective and substantive forms of conflict. Affective conflict corresponds to emotional clashes between individuals or groups, while substantive conflict involves opposition at the intellectual level. Substantive conflict is especially likely to have positive effects on group decisions by promoting better understanding of the issues involved. Affective conflict can also improve group decisions by increasing interest, involvement, and motivation among group members and, in some cases, cohesiveness. On the other hand, affective conflict may cause significant ill will, reduced cohesiveness, and withdrawal by some members from the group process. Baron (1998) provides an interesting discussion of violent conflict, and how it is related to polarized beliefs, group loyalty, and other biases.

Defection and the formation of coalitions is a commonly observed effect of conflict, or power struggles, within groups. Coalitions often form when certain members of the group can gain by following a common course of action at the expense of the long-run objectives of the group as a whole. Rapidly changing coalitions between politicians and political parties are obviously a fact of life. Another typical example is when a subgroup of technical employees leave a corporation to form their own small company producing a product similar to one they had been working on. Coalitions, and their formation, have been examined from decision-analytic and game theory perspectives (Bolton and Chatterjee 1996; Raiffa 1982). These approaches make predictions regarding what coalitions will form, depending on whether the parties are cooperating or competing, which have been tested in a variety of experiments (Bolton and Chatterjee 1996). These experiments have revealed that the formation of coalitions is influenced by expected payoffs, equity issues, and the ease of communication. However, Bazerman (1998) notes that the availability heuristic, overconfidence, and sunk cost effects are likely to explain how coalitions actually form in the real world.

6.2.2. Conflict Resolution

Groups resolve conflict in many different ways. Discussion and argument, voting, negotiation, arbitration, and other forms of third-party intervention are all methods of resolving disputes. Discussion and argument is clearly the most common method followed within groups to resolve conflict. Other methods of conflict resolution normally play a complementary, rather than primary, role in the decision process. That is, the latter methods are relied upon when groups fail to reach consensus after discussion and argument, or they simply serve as the final step in the process.

Group discussion and argument is often viewed as being a less than rational process. Along these lines, Brashers et al. (1994) state that the literature suggests "that argument in groups is a social activity, constructed and maintained in interaction, and guided perhaps by different rules and norms than those that govern the practice of ideal or rational argument. Subgroups speaking with a single voice appear to be a significant force. . . . Displays of support, repetitive agreement, and persistence all appear to function as influence mechanisms in consort with, or perhaps in place of, the quality or rationality of the arguments offered." Brashers et al. also suggest that members of groups appear uncritical because their arguments tend to be consistent with social norms rather than the rules of logic: "[S]ocial rules such as: (a) submission to higher status individuals, (b) experts' opinions are accepted as facts on all matters, (c) the majority should be allowed to rule, (d) conflict and confrontation are to be avoided whenever possible."

A number of approaches for conflict management have been suggested that attempt to address many of the issues raised by Brashers et al. These approaches include seeking consensus rather than allowing decisions to be posed as win–lose propositions, encouraging and training group members to be supportive listeners, deemphasizing status, depersonalizing decision making, and using facilitators (Likert and Likert 1976). Other approaches that have been proposed include directing discussion toward clarifying the issues, promoting an open and positive climate for discussion, facilitating face-saving communications, and promoting the development of common goals (Ellis and Fisher 1994).

Conflicts can also be resolved through voting and negotiation, as discussed further in Section 6.3. Negotiation becomes especially appropriate when the involved people have competing goals and some form of compromise is required. A typical example would be a dispute over pay between a labor union and management. Strategic concerns play a major role in negotiation and bargaining (Schelling 1960). Self-interest on the part of the involved parties is the driving force throughout a process involving threats and promises, proposals and counterproposals, and attempts to discern how the opposing party will respond. Threats and promises are a means of signaling what the response will be to actions taken by an opponent and consequently become rational elements of a decision strategy

(Raiffa 1982). Establishing the credibility of signals sent to an opponent becomes important because if they are not believed, they will not have any influence.

Methods of attaining credibility include establishing a reputation, the use of contracts, cutting off communication, burning bridges, leaving an outcome beyond control, moving in small steps, and using negotiating agents (Dixit and Nalebuff 1991). Given the fundamentally adversarial nature of negotiation, conflict may move from a substantive basis to an affective, highly emotional state. At this stage, arbitration and other forms of third-party intervention may become appropriate due to a corresponding tendency for the negotiating parties to take extreme, inflexible positions.

6.3. Group Performance and Biases

The quality of the decisions made by groups in a variety of different settings has been seriously questioned. Part of the issue here is the phenomenon of so-called group think, which has been blamed for several disastrous public policy decisions (Hart et al. 1997; Janus 1972). Eight symptoms of groupthink cited by Janis and Mann (1977) are the illusion of invulnerability; rationalization (discounting of warnings and negative feedback); belief in the inherent morality of the group; stereotyping of outsiders; pressure on dissenters within the group; self-censorship; illusion of unanimity; and the presence of mindguards who shield the group from negative information. Janis and Mann proposed that the results of groupthink include failure to consider all the objectives and alternatives, failure to reexamine choices and rejected alternatives, incomplete or poor search for information, failure to adequately consider negative information, and failure to develop contingency plans. Groupthink is one of the most cited characteristics of how group decision processes can go wrong. Given the prominence of groupthink as an explanation of group behavior, it is somewhat surprising that only a few studies have empirically evaluated this theory. Empirical evaluation of the groupthink effect and the development of alternative modeling approaches continue to be an active area of research (Hart et al. 1997).

Other research has attempted to measure the quality of group decisions in the real world against rational, or normative, standards. Viscusi (1991) cites several examples of apparent regulatory complacency and regulatory excess in government safety standards in the United States. He also discusses a variety of inconsistencies in the amounts awarded in product liability cases. Baron (1998) provides a long list of what he views as errors in public decision making and their very serious effects on society. These examples include collective decisions resulting in the destruction of natural resources and overpopulation, strong opposition to useful products such as vaccines, violent conflict between groups, and overzealous regulations, such as the Delaney clause. He attributes these problems to commonly held, and at first glance innocuous, intuitions such as Do no harm, Nature knows best, and Be loyal to your own group, the need for retribution (an eye for an eye), and a desire for fairness.

A significant amount of laboratory research also is available that compares the performance of groups to that of individual decision makers (Davis 1992; Kerr et al. 1996). Much of the early work showed that groups were better than individuals on some tasks. Later research indicated that group performance is less than the sum of its parts. Groups tend to be better than individuals on tasks where the solution is obvious once it is advocated by a single member of the group (Davis 1992; Kerr et al. 1996). Another commonly cited finding is that groups tend to be more willing to select risky alternatives than individuals, but in some cases the opposite is true. One explanation is that group interactions cause people within the group to adopt more polarized opinions (Moscovici 1976). Large groups seem especially likely to reach polarized, or extreme, conclusions (Isenberg 1986). Groups also tend to overemphasize the common knowledge of members, at the expense of underemphasizing the unique knowledge certain members have (Gruenfeld et al. 1996; Stasser and Titus, 1985). A more recent finding indicates that groups were more rational than individuals when playing the ultimatum game (Bornstein and Yaniv 1998).

Duffy (1993) notes that teams can be viewed as information processes and cites team biases and errors that can be related to information-processing limitations and the use of heuristics, such as framing. Topics such as mediation and negotiation, jury decision making, and public policy are now being evaluated from the latter perspective (Heath et al. 1994). Much of this research has focused on whether groups use the same types of heuristics and are subject to the same biases of individuals. This research has shown: (1) framing effects and preference reversals (Paese et al. 1993), (2) overconfidence (Sniezek 1992), (3) use of heuristics in negotiation (Bazerman and Neale 1983), and (4) increased performance with cognitive feedback (Harmon and Rohrbaugh 1990). One study indicated that biasing effects of the representativeness heuristic were greater for groups than for individuals (Argote et al. 1986). The conclusion is that group decisions may be better than those of individuals in some situations but are subject to many of the same problems.

6.4. Prescriptive Approaches

A wide variety of prescriptive approaches have been proposed for improving group decision making. The approaches address some of the above issues, including the use of agendas and rules of order, idea-generating techniques such as brainstorming, nominal group and Delphi techniques, decision structuring, and methods of computer-mediated decision making. As noted by Ellis and Fisher (1994), there is conflicting evidence regarding the effectiveness of such approaches. On the negative side, prescriptive approaches might stifle creativity in some situations and can be sabotaged by dissenting members of groups. On the positive side, prescriptive approaches make the decision process more orderly and efficient, promote rational analysis and participation by all members of the group, and help ensure implementation of group decisions. The following discussion briefly reviews some of these tools for improving group decision making.

6.4.1. Agendas and Rules of Order

Agendas and rules of order are often essential to the orderly functioning of groups. As noted by Welch (1994), an agenda ‘‘conveys information about the structure of a meeting: time, place, persons involved, topics to be addressed, perhaps suggestions about background material or preparatory work.’’ Agendas are especially important when the members of a group are loosely coupled or do not have common expectations. Without an agenda, group meetings are likely to dissolve into chaos (Welch 1994). Rules of order, such as *Robert’s Rules of Order* (Robert 1990), play a similarly important role, by regulating the conduct of groups to ensure fair participation, by all group members, including absentees. Rules of order also specify voting rules and means of determining consensus. Decision rules may require unanimity, plurality, or majority vote for an alternative.

Attaining consensus poses an advantage over voting, because voting encourages the development of coalitions, by posing the decision as a win–lose proposition (Ellis and Fisher 1994). Members of the group who voted against an alternative are often unlikely to support it. Voting procedures can also play an important role (Davis 1992).

6.4.2. Idea-Generation Techniques

A variety of approaches have been developed for improving the creativity of groups in the early stages of decision making. Brainstorming is a popular technique for quickly generating ideas (Osborn 1937). In this approach, a small group (of no more than 10 individuals) is given a problem to solve. The members are asked to generate as many ideas as possible. Members are told that no idea is too wild and encouraged to build upon the ideas submitted by others. No evaluation or criticism of the ideas is allowed until after the brainstorming session is finished. Buzz group analysis is a similar approach, more appropriate for large groups (Elliss and Fisher 1994). Here, a large group is first divided into small groups of four to six members. Each small group goes through a brainstorming-like process to generate ideas. They then present their best ideas to the entire group for discussion. Other commonly applied idea-generating techniques include focus group analysis and group exercises intended to inspire creative thinking through role playing (Elliss and Fisher 1994; Clemen 1996).

The use of brainstorming and the other idea-generating methods mentioned above will normally provide a substantial amount of, in some cases, creative suggestions, especially when participants build upon each other’s ideas. However, personality factors and group dynamics can also lead to undesirable results. Simply put, some people are much more willing than others to participate in such exercises. Group discussions consequently tend to center around the ideas put forth by certain more forceful individuals. Group norms, such as deferring to participants with higher status and power, may also lead to undue emphasis on the opinions of certain members.

6.4.3. Nominal Group and Delphi Technique

Nominal group technique (NGT) and the Delphi technique attempt to alleviate some of the disadvantages of working in groups (Delbecq et al. 1975). The nominal group technique consists of asking each member of a group to write down and think about his or her ideas independently. A group moderator then asks each member to present one or more of his or her ideas. Once all of the ideas have been posted, the moderator allows discussion to begin. After the discussion is finished, each participant rates or ranks the presented ideas. The subject ratings are then used to develop a score for each idea. Nominal group technique is intended to increase participation by group members and is based on the idea that people will be more comfortable presenting their ideas if they have a chance to think about them first (Delbecq et al. 1975).

The Delphi technique allows participants to comment anonymously, at their leisure, on proposals made by other group members. Normally, the participants do not know who proposed the ideas they are commenting on. The first step is to send an open-ended questionnaire to members of the group. The results are then used to generate a series of follow-up questionnaires in which more specific questions are asked. The anonymous nature of the Delphi process theoretically reduces the effect of participant status and power. Separating the participants also increases the chance that members will provide opinions ‘‘uncontaminated’’ by the opinions of others.

6.4.4. Structuring Group Decisions

As discussed earlier in this chapter, the field of decision analysis has devised several methods for organizing or structuring the decision-making process. The rational reflection model (Siebold 1992) is a less formal, six-step procedure that serves a similar function. Group members are asked first to define and limit the problem by identifying goals, available resources, and procedural constraints. After defining and limiting the problem, the group is asked to analyze the problem, collect relevant information, and establish the criteria a solution must meet. Potential solutions are then discussed in terms of the agreed-upon decision criteria. After further discussion, the group selects a solution and determines how it should be implemented. The focus of this approach is on forcing the group to confine its discussion to the issues that arise at each step in the decision making process. As such, this method is similar to specifying an agenda.

Raiffa (1982) provides a somewhat more formal decision-analytic approach for structuring negotiations. The approach begins by assessing (1) the alternatives to a negotiated settlement, (2) the interests of the involved parties, and (3) the relative importance of each issue. This assessment allows the negotiators to think analytically about mutually acceptable solutions. In certain cases, a bargaining zone is available. For example, an employer may be willing to pay more than the minimum salary acceptable to a potential employee. In this case, the bargaining zone is the difference between the maximum salary the employer is willing to pay and the minimum salary a potential employee is willing to accept. The negotiator may also think about means of expanding the available resources to be divided, potential trading issues, or new options that satisfy the interests of the concerned parties.

Other methods for structuring group preferences are discussed in Keeny and Raiffa (1976). The development of group utility functions is one such approach. A variety of computer-mediated methods for structuring group decisions are also available, as discussed below.

6.4.5. Computer-Mediated Group Decision Making

Computer tools for helping groups make decisions are now available that implement all of the approaches discussed above to varying degrees. The spectrum of available tools ranges from traditional tools used in decision analysis, such as the analytic hierarchy process (Basak and Saaty 1993; Saaty 1988), to electronic meeting places (Mockler and Dologite 1991; Nunamaker et al. 1991) and group decision support systems (GDSS) (Sage 1997). Some of the many functions provided by GDSS (Johansen 1988; Sage 1997) include computer-supported presentations, project and calendar management, group authoring, electronic meeting places, audio and visual conferencing, screen sharing, group e-mail services, text-filtering services, conversational structuring, group and organizational memory management, and comprehensive work team support.

Computer-mediated group decision making has several potential benefits (Brasher et al. 1994), including (1) enabling all participants to work simultaneously (they don't have to wait their turn to speak), (2) providing a more equal and potentially anonymous opportunity to be heard, (3) providing a more structured environment (that is, a more linear process and control of the agenda). Computer-mediated group decision making also might make it easier to control and manage conflict, through the use of facilitators and convenient voting procedures. As noted by Sage (1997), the purpose of GDSS is to (1) remove communication barriers, (2) provide techniques for the formulation, analysis, and interpretation of decisions, and (3) systematically direct the decision-making process. Successfully attaining these objectives means that a GDSS must provide the right information to decision makers, at an appropriate level of detail, at the time it is needed, in a form that is conveniently applied. Simply put, to be useful, a GDSS must not make group decision making more difficult than it already is.

7. SUMMARY CONCLUSIONS

Beach (1993) discusses four revolutions in behavioral decision theory. The first took place when it was recognized that the evaluation of alternatives is seldom extensive. It is illustrated by use of the satisficing rule (Simon 1955) and heuristics (Tversky and Kahneman 1974) rather than optimizing. The second occurred when it was recognized that people choose between strategies to make decisions. It is marked by the development of contingency theory (Beach 1990) and cognitive continuum theory (Hammond 1980). The third is currently occurring. It involves the realization that people rarely make choices and instead rely on prelearned procedures. This perspective is illustrated by the levels-of-processing approach (Rasmussen 1983) and recognition-primed decisions (Klein 1989). The fourth is just beginning. It involves recognization that decision-making research must abandon a single-minded focus on the economic view of decision making and include approaches drawn from relevant developments and research in cognitive psychology, organizational behavior, and systems theory.

The discussion within this chapter parallels this view of decision making. The integrative model presented at the beginning of the chapter shows how these various approaches fit together as a whole. Each path through the model is distinguished by specific sources of conflict, the methods of conflict resolution followed, and the types of decision rules used to analyze the results of conflict-resolution processes. The different paths through the model correspond to fundamentally different ways of making decisions, ranging from routine situation assessment-driven decisions to satisficing, analysis of single- and multiattribute expected utility, and even obtaining consensus of multiple decision makers in group contexts. Numerous other strategies discussed in this chapter are also described by particular paths through the model.

This chapter goes beyond simply describing methods of decision making by pointing out reasons people and groups may have difficulty making good decisions. These include cognitive limitations, inadequacies of various heuristics used, biases and inadequate knowledge of decision makers, and task-related factors, such as risk, time pressure, and stress. The discussion also provides insight into the effectiveness of approaches for improving human decision making. The models of selective attention point to the value of providing only truly relevant information to decision makers. Irrelevant information might be considered simply because it is there, especially if it is highly salient. Methods of highlighting or emphasizing relevant information, therefore, clearly seem to be warranted. The models of selective information also indicate that methods of helping decision makers cope with working memory limitations will be of value. There also is reason to believe that providing feedback to decision makers in dynamic decision-making situations will be useful. Cognitive, rather than outcome, feedback is indicated as being particularly helpful when decision makers are learning. Training decision makers also seems to offer potentially large benefits. One reason for this conclusion is that the studies of naturalistic decision making revealed that most decisions are made on a routine, nonanalytical basis.

The studies of debiasing also partially support the potential benefits of training and feedback. On the other hand, the many failures to debias expert decision makers imply that decision aids, methods of persuasion, and other approaches intended to improve decision making are no panacea. Part of the problem is that people tend to start with preconceived notions about what they should do and show a tendency to seek out and bolster confirming evidence. Consequently, people show a tendency to develop overconfidence with experience, and strongly held beliefs become difficult to modify, even if they are hard to defend rationally.

REFERENCES

Ainslie, G. (1975), "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin,* Vol. 82, pp. 463–509.

Allais, M. (1953), "Le comportement de l'homme rationel devant le risque: Critique des postulates et axioms de l'ecole americaine," *Econometrica,* Vol. 21, pp. 503–546.

Anderson, N. H. (1981), *Foundations of Information Integration Theory,* Academic Press, New York.

Argote, L., Seabright, M. A., and Dyer, L. (1986), "Individual versus Group: Use of Base-Rate and Individuating Information," *Organizational Behavior and Human Decision Making Processes,* Vol. 38, pp. 65–75.

Arkes, H. R., and Blumer, C. (1985), "The Psychology of Sunk Cost," *Organizational Behavior and Human Decision Processes,* Vol. 35, pp. 124–140.

Balzer, W. K., Doherty, M. E., and O'Connor, R. O., Jr. (1989), "Effects of Cognitive Feedback on Performance," *Psychological Bulletin,* Vol. 106, pp. 41–433.

Bar-Hillel, M. (1973), "On the Subjective Probability of Compound Events," *Organizational Behavior and Human Performance,* Vol. 9, pp. 396–406.

Baron, J. (1985), *Rationality and Intelligence,* Cambridge University Press, Cambridge.

Baron, J. (1998), *Judgment Misguided: Intuition and Error in Public Decision Making,* Oxford University Press, New York.

Basak, I., and Saaty, T. (1993), "Group Decision Making Using the Analytic Hierarchy Process," *Mathematical and Computer Modeling,* Vol. 17, pp. 101–109.

Bazermen, M. (1998), *Judgment in Managerial Decision Making,* 4th Ed., Wiley & Sons, New York.

Bazerman, M. H., and Neale, M. A. (1983), "Heuristics in Negotiation: Limitations to Effective Dispute Resolution," in *Negotiating in Organizations,* M. H. Bazerman and R. Lewicki, Eds., Sage, Beverly Hills.

Beach, L. R. (1990), *Image Theory: Decision Making in Personal and Organizational Contexts,* John Wiley & Sons, Chichester.

Beach, L. R. (1993), "Four Revolutions in Behavioral Decision Theory," in *Leadership Theory and Research,* M. M. Chemers and R. Ayman, Eds., Academic Press, San Diego.

Beach, L. R., and Mitchell, T. R. (1978), "A Contingency Model for the Selection of Decision Strategies," *Academy of Management Journal,* Vol. 3, pp. 439–449.

Bell, D. (1982), "Regret in Decision Making under Uncertainty," *Operations Research,* Vol. 30, pp. 961–981.

Bernoulli, D. (1738), *Exposition of a New Theory of the Measurement of Risk,* Imperial Academy of Science, St. Petersburg.

Birnbaum, M. H., and Mellers, B. A. (1983), "Bayesian Inference: Combining Base Rates with Opinions of Sources Who Vary in Credibility," *Journal of Personality and Social Psychology,* Vol. 37, pp. 792–804.

Birnbaum, M. H., Coffey, G., Mellers, B. A., and Weiss, R. (1992), "Utility Measurement: Configural-Weight Theory and the Judge's Point of View," *Journal of Experimental Psychology: Human Perception and Performance,* Vol. 18, pp. 331–346.

Bock, R. D., and Jones, L. V. (1968), *The Measurement and Prediction of Judgment and Choice,* Holden-Day, San Francisco.

Bolton, G. E., and Chatterjee, K. (1996), "Coalition Formation, Communication, and Coordination: An Exploratory Experiment," in *Wise Choices: Decisions, Games, and Negotiations,* R. J. Zeckhauser, R. L. Keeney, and J. K. Sebenius, Eds., Harvard Business School Press, Boston.

Bornstein, G., and Yaniv, I. (1998), "Individual and Group Behavior in the Ultimatum Game: Are Groups More Rational Players?" *Experimental Economics,* Vol. 1, pp. 101–108.

Brashers, D. E., Adkins, M., and Meyers, R. A. (1994), "Argumentation and Computer-Mediated Group Decision Making," 12 in *Group Communication in Context,* L. R. Frey, Ed., Erlbaum, Hillsdale, NJ.

Brehmer, B. (1981), "Models of Diagnostic Judgment," in *Human Detection and Diagnosis of System Failures,* J. Rasmussen and W. Rouse, Eds., Plenum Press, New York.

Brehmer, B., and Joyce, C. R. B. (1988), *Human Judgment: The SJT View,* North-Holland, Amsterdam.

Brookhouse, J. K., Guion, R. M., and Doherty, M. E. (1986), "Social Desirability Response Bias as One Source of the Discrepancy between Subjective Weights and Regression Weights," *Organizational Behavior and Human Decision Processes,* Vol. 37, pp. 316–328.

Brunswick, E. (1952), *The Conceptual Framework of Psychology,* University of Chicago Press, Chicago.

Buck, J. R. (1989), *Economic Risk Decisions in Engineering and Management,* Iowa State University Press, Ames, IA.

Budescu, D., and Weiss, W. (1987), "Reflection of Transitive and Intransitive Preferences: A Test of Prospect Theory," *Organizational Behavior and Human Performance,* Vol. 39, pp. 184–202.

Busemeyer, J. R., and Townsend, J. T. (1993), "Decision Field Theory: A Dynamic–Cognitive Approach to Decision Making in an Uncertain Environment," *Psychological Review,* Vol. 100, pp. 432–459.

Caverni, J. P., Fabre, J. M., and Gonzalez, M. (1990), *Cognitive Biases,* North-Holland, Amsterdam.

Chemers, M. M., and Ayman, R. (eds) (1993), *Leadership Theory and Research,* Academic Press, San Diego.

Cheng, P. E., and Holyoak, K. J. (1985), "Pragmatic Reasoning Schemas," *Cognitive Psychology,* Vol. 17, pp. 391–416.

Christensen-Szalanski, J. J., and Willham, C. F. (1991), "The Hindsight Bias: A Meta-Analysis," *Organizational Behavior and Human Decision Processes,* Vol. 48, pp. 147–168.

Clemen, R. T. (1996), *Making Hard Decisions: An Introduction to Decision Analysis,* 2nd Ed., Duxbury Press, Belmont, CA.

Cohen, M. S. (1993), "The Naturalistic Basis of Decision Biases," in *Decision Making in Action: Models and Methods,* G. A. Klein, J. Orasanu, R. Calderwood, and E. Zsambok, Eds., Ablex, Norwood, NJ, pp. 51–99.

Connolly, T., Ordonez, L. D., and Coughlan, R. (1997), "Regret and Responsibility in the Evaluation of Decision Outcomes," *Organizational Behavior and Human Decision Processes,* Vol. 70, pp. 73–85.

Davis, J. H. (1992), "Some Compelling Intuitions about Group Consensus Decisions, Theoretical and Empirical Research, and Interperson Aggregation Phenomena: Selected Examples, 1950–1990," *Organizational Behavior and Human Decision Processes,* Vol. 52, pp. 3–38.

Dawes, R. M., and Mulford, M. (1996), "The False Consensus Effect and Overconfidence: Flaws in Judgement or Flaws in How We Study Judgement?" *Organizational Behavior and Human Decision Processes,* Vol. 65, pp. 201–211.

Dawes, R. M., van de Kragt, A. J. C., and Orbell, J. M. (1988), "Not Me or Thee but We: The Importance of Group Identity in Eliciting Cooperation in Dilemma Situations: Experimental Manipulations," *Acta Psychologia,* Vol. 68, pp. 83–97.

Delbecq, A. L., Van de Ven, A. H., and Gustafson, D. H. (1975), *Group Techniques for Program Planning,* Scott, Foresman, Glenview, IL.

Dixit, A., and Nalebuff, B. (1991), "Making Strategies Credible," in *Strategy and Choice,* R. J. Zechhauser, Ed., MIT Press, Cambridge, MA, pp. 161–184.

Dorris, A. L., and Tabrizi, J. L. (1978), "An Empirical Investigation of Consumer Perception of Product Safety," *Journal of Products Liability,* Vol. 2, pp. 155–163.

Driskell, J. E., and Salas, E., Eds. (1996), *Stress and Human Performance,* Erlbaum, Hillsdale, NJ.

Drucker, P. F. (1985), *The Effective Executive,* Harper Row, New York.

Du Charme, W. (1970), "Response Bias Explanation of Conservative Human Inference," *Journal of Experimental Psychology,* Vol. 85, pp. 66–74.

Duda, R. O., Hart, K., Konolige, K., and Reboh, R. (1979), "A Computer-Based Consultant for Mineral Exploration," Technical Report, SRI International, Stanford, CA.

Duffy, L. (1993), "Team Decision Making Biases: An Information Processing Perspective," in *Decision Making in Action: Models and Methods,* G. A. Klein, J. Orasanu, R. Calderwood, and E. Zsambok, Eds., Ablex, Norwood, NJ.

Edland, E., and Svenson, O. (1993), "Judgment and Decision Making under Time Pressure," in *Time Pressure and Stress in Human Judgment and Decision Making,* O. Svenson, and A. J. Maule, Eds., Plenum, New York.

Edwards, W. (1954), "The Theory of Decision Making," *Psychological Bulletin,* Vol. 41, pp. 380–417.

Edwards, W. (1968), "Conservatism in Human Information Processing," in *Formal Representation of Human Judgment,* B. Kleinmuntz, Ed., John Wiley & Sons, New York, pp. 17–52.

Einhorn, H. J., and Hogarth, R. M. (1978), "Confidence in Judgment: Persistence of the Illusion of Validity," *Psychological Review,* Vol. 70, pp. 193–242.

Einhorn, H. J., and Hogarth, R. M. (1981), "Behavioral Decision Theory: Processes of Judgment and Choice," *Annual Review of Psychology,* Vol. 32, pp. 53–88.

Ellsberg, D. (1961), "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics,* Vol. 75, pp. 643–699.

Elster, J., Ed. (1986), *The Multiple Self,* Cambridge University Press, Cambridge.

Embrey, D. E. (1984), *SLIM-MAUD: An Approach to Assessing Human Error Probabilities Using Structured Expert Judgment,* NUREG/CR-3518, Vols. 1 and 2, U.S. Nuclear Regulatory Commission, Washington, DC.

Ericsson, K. A., and Simon, H. A. (1984), *Protocol Analysis: Verbal Reports as Data,* MIT Press, Cambridge, MA.

Etzioni, A. (1988), "Normative-Affective Factors: Toward a New Decision-Making Model," *Journal of Economic Psychology,* Vol. 9, pp. 125–150.

Evans, J. B. T. (1989), *Bias in Human Reasoning: Causes and Consequences,* Erlbaum, London.

Evans, J. B. T., and Pollard, P. (1985), "Intuitive Statistical Inferences about Normally Distributed Data," *Acta Psychologica,* Vol. 60, pp. 57–71.

Feather, N. T. (1966), "Effects of Prior Success and Failure on Expectations of Success and Failure," *Journal of Personality and Social Psychology,* Vol. 3, pp. 287–298.

Fedrizzi, M., Kacprzyk, J., and Yager, R. R., Eds. (1994), *Decision Making under Dempster–Shafer Uncertainties,* John Wiley & Sons, New York.

Fischhoff, B. (1982), "For Those Condemned to Study the Past: Heuristics and Biases in Hindsight," in *Judgment under Uncertainty: Heuristics and Biases,* D. Kahneman, P. Slovic, and A. Tversky, Eds., Cambridge University Press, Cambridge.

Fischhoff, B., and MacGregor, D. (1982), "Subjective Confidence in Forecasts," *Journal of Forecasting,* Vol. 1, pp. 155–172.

Fischhoff, B., Slovic, P., and Lichtenstein, S. (1977), "Knowing with Certainty: The Appropriateness of Extreme Confidence," *Journal of Experimental Psychology: Human Perception and Performance,* Vol. 3, pp. 552–564.

Fischhoff, B., Slovic, P., and Lichtenstein, S. (1978), "Fault Trees: Sensitivity of Estimated Failure Probabilities to Problem Representation," *Journal of Experimental Psychology: Human Perception and Performance,* Vol. 4, pp. 330–344.

Fishburn, P. C. (1974), "Lexicographic Orders, Utilities, and Decision Rules: A Survey," *Management Science,* Vol. 20, pp. 1442–1471.

Flin R., Salas, E., Strub, M., and Martin, L., Eds. (1997), *Decision Making under Stress: Emerging Themes and Applications,* Ashgate, Aldershot, UK.

Fong, G. T., Krantz, D. H., and Nisbett, R. E. (1986), "The Effects of Statistical Training on Thinking about Everyday Problems," *Cognitive Psychology,* Vol. 18, pp. 253–292.

Friedman, J. W. (1990), *Game Theory with Applications to Economics,* Oxford University Press, New York.

Frisch, D., and Clemen, R. T. (1994), "*Beyond Expected Utility: Rethinking Behavioral Decision Research,*" *Psychological Bulletin,* Vol. 116, No. 1, pp. 46–54.

Gertman, D. I., and Blackman, H. S. (1994), *Human Reliability and Safety Analysis Data Handbook,* John Wiley & Sons, New York.

Goldstein, W. M., and Hogarth, R. M. (1997), "Judgment and Decision Research: Some Historical Context," in *Research on Judgment and Decision Making: Currents, Connections, and Controversies,* W. M. Goldstein, and R. M. Hogarth, Eds., Cambridge University Press, Cambridge, pp. 3–65.

Gordon, J., and Shortliffe, E. H. (1984), "The Dempster–Schafer Theory of Evidence," in *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project,* Addison-Wesley, Reading, MA.

Gruenfeld, D. H., Mannix, E. A., Williams, K. Y., and Neale, M. A. (1996), "Group Composition and Decision Making: How Member Familiarity and Information Distribution Affect Process and Performance," *Organizational Behavior and Human Decision Making Processes,* Vol. 67:1, pp. 1–15.

Güth, W., Schmittberger, R., and Schwarze, B. (1982), "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization,* Vol. 3, pp. 367–388.

Hamilton, V., and Warburton, D. M., Eds. (1979), *Human Stress and Cognition,* John Wiley & Sons, New York.

Hammer, W. (1993), *Product Safety Management and Engineering,* 2nd Ed., ASSE, Chicago.

Hammond, K. R. (1980), "Introduction to Brunswikian Theory and Methods," in *Realizations of Brunswick's Experimental Design,* K. R. Hammond, and N. E. Wascoe, Eds., Jossey-Bass, San Francisco.

Hammond, K. R. (1993), "Naturalistic Decision Making From a Brunswikian Viewpoint: Its Past, Present, Future," in *Decision Making in Action: Models and Methods,* G. A. Klein, J. Orasanu, R. Calderwood, and E. Zsambok, Eds., Ablex, Norwood, NJ, pp. 205–227.

Hammond, K. R, Stewart, T. R., Brehmer, B., and Steinmann, D. O. (1975), "Social Judgment Theory," in *Human Judgment and Decision Processes,* M. F. Kaplan, and S. Schwartz, Eds., Academic Press, New York, pp. 271–312.

Hammond, K. R., Hamm, R. M., Grassia, J., and Pearson, T. (1987), "Direct Comparison of the Efficacy of Intuitive and Analytical Cognition in Expert Judgment," *IEEE Transactions on Systems, Man, and Cybernetics,* Vol. SMC-17, pp. 753–770.

Hardin, G. (1968), "The Tragedy of the Commons," *Science,* Vol. 162, pp. 1243–1248.

Harmon, J., and Rohrbaugh, J. (1990), "Social Judgement Analysis and Small Group Decision Making: Cognitive Feedback Effects on Individual and Collective Performance," *Organizational Behavior and Human Decision Processes,* Vol. 46, pp. 34–54.

Hart, P., Stern, E. K., and Sundelius, B. (1997), *Beyond Groupthink: Political Group Dynamics and Foreign Policy-Making,* University of Michigan Press, Ann Arbor.

Heath, C. (1995), "Escalation and De-escalation of Commitment in Response to Sunk Costs: The Role of Budgeting in Mental Accounting," *Organizational Behavior and Human Decision Processes,* Vol. 62, pp. 38–54.

Heath, L., Tindale, R. S., Edwards, J., Posavac, E. J., Bryant, F. B., Henderson-King, E., Suarez-Balcazar, Y., and Myers, J. (1994), *Applications of Heuristics and Biases to Social Issues,* Plenum Press, New York.

Hogarth, R. M., and Einhorn, H. J. (1992), "Order Effects in Belief Updating: The Belief-Adjustment Model," *Cognitive Psychology,* Vol. 24, pp. 1–55.

Holtzman, S. (1989), *Intelligent Decision Systems,* Addison-Wesley, Reading, MA.

Howard, R. A. (1968), "The Foundations of Decision Analysis," *IEEE Transactions on Systems, Science, and Cybernetics,* Vol. SSC-4, pp. 211–219.

Howard, R. A. (1988), "Decision Analysis: Practice and Promise," *Management Science,* Vol. 34, pp. 679–695.

Huber, J., Wittink, D. R., Fiedler, J. A., and Miller, R. (1993), "The Effectiveness of Alternative Preference Elicitation Procedures in Predicting Choice," *Journal of Marketing Research,* Vol. 30, pp. 105–114.

Isen, A. M. (1993), "Positive Affect and Decision Making," in *Handbook of Emotions,* M. Lewis, and J. M. Haviland, Eds., Guilford Press, New York, pp. 261–277.

Isenberg, D. J. (1986), "Group Polarization: A Critical Review and Meta Analysis," *Journal of Personality and Social Psychology,* Vol. 50, pp. 1141–1151.

Iverson, G., and Luce, R. D. (1998), "The Representational Measurement Approach to Psychophysical and Judgmental Problems," in *Measurement, Judgement, and Decision Making,* M. H. Birnbaum, Ed., Academic Press, San Diego, pp. 1–79.

Jacoby, J. (1977), "Information Load and Decision Quality: Some Contested Issues," *Journal of Marketing Research,* Vol. 14, pp. 569–573.

Janis, I. L. (1972), *Victims of Groupthink,* Houghton-Mifflin, Boston.

Janis, I. L., and Mann, L. (1977), *Decision Making: A Psychological Analysis of Conflict, Choice, and Commitment,* Free Press, New York.

Johansen, R. (1988), *Groupware: Computer Support for Business Teams,* Free Press, New York.

Johnson, E. J., and Payne, J. W. (1985), "Effort and Accuracy in Choice," *Management Science,* Vol. 31, No. 4, pp. 395–414.

Kahneman, D., and Tversky, A. (1973), "On the Psychology of Prediction," *Psychological Review,* Vol. 80, pp. 251–273.

Kahneman, D., and Tversky, A. (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica,* Vol. 47, pp. 263–291.

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

Keeney, R. L., and Raiffa, H. (1976), *Decisions with Multiple Objectives: Preferences and Value Tradeoffs,* John Wiley & Sons, New York.

Kerr, L. N., MacCoun, R. J., and Kramer, G. P. (1996), "Bias in Judgment: Comparing Individuals and Groups," *Psychological Review,* Vol. 103, pp. 687–719.

Kieras, D. E. (1985), "The Role of Prior Knowledge in Operating Equipment from Written Instructions," Report No. 19 (FR-85/ONR-19), Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor.

Keren, G. (1990), "Cognitive Aids and Debiasing Methods: Can Cognitive Pills Cure Cognitive Ills?" in *Cognitive Biases,* J. P. Caverni, J. M. Fabre, and M. Gonzalez, Eds., Amsterdam, North Holland.

Klayman, J., and Ha, Y. W. (1987), "Confirmation, Disconfirmation, and Information in Hypothesis Testing," *Journal of Experimental Psychology: Human Learning and Memory,* pp. 211–228.

Klein, G. A. (1989), "Recognition-Primed Decisions," in Advances in Man-Machine System Research, W. Rouse, Ed., CT: JAI Press, Greenwich.

Klein, G. A. (1995), "The Value Added by Cognitive Analysis," in *Proceedings of the Human Factors and Ergonomics Society—39th Annual Meeting,* pp. 530–533.

Klein, G. A. (1996), "The Effect of Acute Stressors on Decision Making," in *Stress and Human Performance,* J. E. Driskell, and E. Salas, Eds., Erlbaum, Hillsdale, NJ.

Klein, G. A. (1998), *Sources of Power: How People Make Decisions,* MIT Press, Cambridge, MA.

Klein, G. A., and Wolf, S. (1995), "Decision-Centered Training," *Proceedings of the Human Factors and Ergonomics Society—39th Annual Meeting,* pp. 1249–1252.

Klein, G. A., Orasanu, J., Calderwood, R., and Zsambok, E., Eds. (1993), *Decision Making in Action: Models and Methods,* Ablex, Norwood, NJ.

Kleinmuntz, B. (1984), "The Scientific Study of Clinical Judgment in Psychology and Medicine," *Clinical Psychology Review,* Vol. 4, pp. 111–126.

Koehler, J. J. (1996), "The Base Rate Fallacy Reconsidered: Descriptive, Normative, and Methodological Challenges," *Behavioral and Brain Sciences,* Vol. 19, pp. 1–53.

Koriat, A., Lichtenstein, S., and Fischhoff, B. (1980), "Reasons for Confidence," *Journal of Experimental Psychology: Human Learning and Memory,* Vol. 6, pp. 107–118.

Kraus, N. N., and Slovic, P. (1988), "Taxonomic Analysis of Perceived Risk: Modeling Individual and Group Perceptions within Homogenous Hazards Domains," *Risk Analysis,* Vol. 8, pp. 435–455.

Kuhberger, A. (1995), "The Framing of Decisions: A New Look at Old Problems," *Organizational Behavior and Human Decision Processes,* Vol. 62, pp. 230–240.

Lathrop, R. G. (1967), "Perceived Variability," *Journal of Experimental Psychology,* Vol. 23, pp. 498–502.

Lehto, M. R. (1991), "A Proposed Conceptual Model of Human Behavior and Its Implications for Design of Warnings," *Perceptual and Motor Skills,* Vol. 73, pp. 595–611.

Lehto, M. R., and Papastavrou, J. (1991), "A Distributed Signal Detection Theory Model: Implications to the Design of Warnings," in *Proceedings of the 1991 Automatic Control Conference* (Boston), pp. 2586–2590.

Lehto, M. R., Papastavrou, J. P., Ranney, T. A., and Simmons, L. (2000), "An Experimental Comparison of Conservative vs Optimal Collision Avoidance System Thresholds," *Safety Science,* Vol. 36, No. 3, pp. 185–209.

Lehto, M. R., James, D. S., and Foley, J. P. (1994), "Exploratory Factor Analysis of Adolescent Attitudes Toward Alcohol and Risk," *Journal of Safety Research,* Vol. 25, pp. 197–213.

Levin, L. P. (1975), "Information Integration in Numerical Judgements and Decision Processes," *Journal of Experimental Psychology: General,* Vol. 104, pp. 39–53.

Lichtenstein, S., Fischhoff, B., and Phillips, L. D. (1982), "Calibration of Probabilities: The State of the Art to 1980," in *Judgment under Uncertainty: Heuristics and Biases,* D. Kahneman, P. Slovic, and A. Tversky, Eds., pp. 306–334.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., and Coombs, B. (1978), "Judged Frequency of Lethal Events," *Journal of Experimental Psychology: Human Learning & Memory,* Vol. 4, pp. 551–578.

Likert, R., and Likert, J. G. (1976), *New Ways of Managing Conflict,* McGraw-Hill, New York.

Luce, R. D., and Raiffa, H. (1957), *Games and Decisions,* John Wiley & Sons, New York.

Mallet, L., Vaught, C., and Brnich, M. J., Jr. (1993), "Sociotechnical Communication in an Underground Mine Fire: A Study of Warning Messages During an Emergency Evacuation," *Safety Science,* Vol. 16, pp. 709–728.

Martocchio, J. J., Webster, J., and Baker, C. R. (1993), "Decision-Making in Management Information Systems Research: The Utility of Policy Capturing Methodology," *Behaviour and Information Technology,* Vol. 12, pp. 238–248.

Maule, A. J., and Hockey, G. R. J. (1993), "State, Stress, and Time Pressure," in *Time Pressure and Stress in Human Judgment and Decision Making,* O. Svenson, and A. J. Maule, Eds., Plenum, New York, pp. 27–40.

McGuire, W. J. (1966), "Attitudes and Opinions," *Annual Review of Psychology,* Vol. 17, pp. 475–514.

McKnight, A. J., Langston, E. A., McKnight, A. S., and Lange, J. E. (1995), "The Bases of Decisions Leading to Alcohol Impaired Driving," in *Proceedings of the 13th International Conference on Alcohol, Drugs, and Traffic Safety* (Adelaide, August 13–18), C. N. Kloeden and A. J. McLean, Eds., pp. 143–147.

Messick, D. M. (1991), "Equality as a Decision Heuristic," in *Psychological Issues in Distributive Justice,* B. Mellers, Ed., Cambridge University Press, New York.

Mockler, R. L., and Dologite, D. G. (1991), "Using Computer Software to Improve Group Decision-Making," *Long Range Planning,* Vol. 24, pp. 44–57.

Montgomery, H. (1989), "From Cognition to Action: The Search for Dominance in Decision Making," in *Process and Structure in Human Decision Making,* H. Montgomery and O. Svenson, Eds., John Wiley & Sons, Chichester.

Moscovici, S. (1976), *Social Influence and Social Change,* Academic Press, London.

Murphy, A. H., and Winkler, R. L. (1974), "Probability Forecasts: A Survey of National Weather Service Forecasters," *Bulletin of the American Meteorological Society,* Vol. 55, pp. 1449–1453.

Myers, J. L., Suydam, M. M., and Gambino, B. (1965), "Contingent Gains and Losses in a Risky Decision Situation," *Journal of Mathematical Psychology,* Vol. 2, pp. 363–370.

Naruo, N., Lehto, M., and Salvendy, G. (1990), "Development of a Knowledge Based Decision Support System for Diagnosing Malfunctions of Advanced Production Equipment," *International Journal of Production Research,* Vol. 28, pp. 2259–2276.

Navon, D. (1979), "The Importance of Being Conservative," *British Journal of Mathematical and Statistical Psychology,* Vol. 31, pp. 33–48.

Newell, A., and Simon, H. A. (1972), *Human Problem Solving,* Prentice-Hall, Englewood Cliffs, NJ.

Nisbett, R., and Ross, L. (1980), *Human Inference: Strategies and Shortcomings of Social Judgment,* Prentice-Hall, Englewood Cliffs, NJ.

Nunamaker, J. F., Dennis, A. R., Valacich, J. S., Vogel, D. R., and George, J. F. (1991), "Electronic Meeting Systems to Support Group Work: Theory and Practice at Arizona," *Communications of the ACM,* Vol. 34, pp. 40–61.

Orasanu, J. (1997), "Stress and Naturalistic Decision Making: Strengthening the Weak Links," in. *Decision Making under Stress: Emerging Themes and Applications,* R. Flin, E. Salas, M. Strub, and L. Martin, Eds., Ashgate, Aldershot, UK.

Orasanu, J., and Salas, E. (1993), "Team Decision Making in Complex Environments," in *Decision Making in Action: Models and Methods,* G. A. Klein, J. Orasanu, R. Calderwood, and E. Zsambok, Eds., Ablex, Norwood, NJ.

Osborn, F. (1937), *Applied Imagination,* Charles Scribner & Sons, New York.

Paese, P. W., Bieser, M., and Tubbs, M. E. (1993), "Framing Effects and Choice Shifts in Group Decision Making," *Organizational Behavior and Human Decision Processes,* Vol. 56, pp. 149–165.

Payne, J. W. (1980), "Information Processing Theory: Some Concepts and Methods Applied to Decision Research," in *Cognitive Processes in Choice and Decision Research,* T. S. Wallsten, Ed., Erlbaum, Hillsdale, NJ.

Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993), *The Adaptive Decision Maker,* Cambridge University Press, Cambridge.

Pennington, N., and Hastie, R. (1986), "Evidence Evaluation in Complex Decision Making," *Journal of Personality and Social Psychology,* Vol. 51, pp. 242–258.

Pennington, N., and Hastie, R. (1988), "Explanation-Based Decision Making: Effects of Memory Structure on Judgment," *Journal of Experimental Psychology: Learning, Memory, and Cognition,* Vol. 14, pp. 521–533.

Pitz, G. F. (1980), "The Very Guide of Life: The Use of Probabilistic Information for Making Decisions," in *Cognitive Processes in Choice and Decision Behavior,* T. S. Wallsten, Ed., Erlbaum, Hillsdale, NJ.

Raiffa, H. (1968), *Decision Analysis,* Addison-Wesley, Reading, MA.

Raiffa, H. (1982), *The Art and Science of Negotiation,* Harvard University Press, Cambridge, MA.

Rasmussen, J. (1983), "Skills, Rules, Knowledge: Signals, Signs, and Symbols and Other Distinctions in Human Performance Models," *IEEE Transactions on Systems, Man, and Cybernetics,* Vol. SMC-13, No. 3, pp. 257–267.

Reason, J. (1990), *Human Error,* Cambridge University Press, Cambridge.

Rethans, A. J. (1980), "Consumer Perceptions of Hazards," in *PLP-80 Proceedings,* pp. 25–29.

Robert, H. M. (1990), *Robert's Rules of Order Newly Revised,* 9th Ed., Scott, Foresman, Glenview, IL.

Saaty, T. L. (1988), *Multicriteria Decision Making: The Analytic Hierarchy Process,* T. Saaty, Pittsburgh.

Sage, A. (1981), "Behavioral and Organizational Considerations in the Design of Information Systems and Processes for Planning and Decision Support," *IEEE Transactions on Systems, Man, and Cybernetics,* Vol. SMC-11, pp. 61–70.

Savage, L. J. (1954), *The Foundations of Statistics,* John Wiley & Sons, New York.

Schelling, T. (1960), *The Strategy of Conflict,* Harvard University Press, Cambridge, MA.

Schelling, T. (1978), *Micromotives and Macrobehavior,* W. W. Norton, New York.

Selye, H. (1936), "A Syndrome Produced by Diverse Noxious Agents," *Nature,* Vol. 138, p. 32.

Selye, H. (1979), "The Stress Concept and Some of Its Implications," in *Human Stress and Cognition,* V. Hamilton and D. M. Warburton, Eds., John Wiley & Sons, New York.

Shafer, G. (1976), *A Mathematical Theory of Evidence,* Princeton University Press, Princeton, NJ.

Siebold, D. R. (1992), "Making Meetings More Successful: Plans, Formats, and Procedures for Group Problem-Solving," in *Small Group Communication,* 6th Ed., R. Cathcart and L. Samovar, Eds., Brown, Dubuque, IA, pp. 178–191.

Simon, H. A. (1955), "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics,* Vol. 69, pp. 99–118.

Simon, H. A. (1983), "Alternative Visions of Rationality," in *Reason in Human Affairs,* Stanford University Press, Stanford, CA.

Singleton, W. T., and Hovden, J. (1987), *Risk and Decisions,* John Wiley & Sons, New York.

Slovic, P. (1978), "The Psychology of Protective Behavior," *Journal of Safety Research,* Vol. 10, pp. 58–68.

Slovic, P. (1987), "Perception of Risk," *Science,* Vol. 236, pp. 280–285.

Slovic, P., Fischhoff, B., and Lichtenstein, S. (1977), "Behavioral Decision Theory," *Annual Review of Psychology,* Vol. 28, pp. 1–39.

Sniezek, J. A. (1992), "Groups under Uncertainty: An Examination of Confidence in Group Decision Making," *Organizational Behavior and Human Decision Processes,* Vol. 52, pp. 124–155.

Sniezek, J. A., and Henry, R. A. (1989), "Accuracy and Confidence in Group Judgment," *Organizational Behavior and Human Decision Processes,* Vol. 43, pp. 1–28.

Stanoulov, N. (1994), "Expert Knowledge and Computer-Aided Group Decision Making: Some Pragmatic Reflections," *Annals of Operations Research,* Vol. 51, pp. 141–162.

Stasser, G., and Titus, W. (1985), "Pooling of Unshared Information in Group Decision Making: Biased Information Sampling during Discussion," *Journal of Personality and Social Psychology,* Vol. 48, pp. 1467–1478.

Stevenson, M. K., Busemeyer, J. R., and Naylor, J. C. (1993), "Judgment and Decision-Making Theory," in *Handbook of Industrial and Organizational Psychology,* 2nd Ed., Vol. 1, M. D. Dunnette and L. M. Hough, Eds., Consulting Psychologists Press, Palo Alto, CA.

Stone, M. (1960), "Models for Reaction Time," *Psychometrika,* Vol. 25, pp. 251–260.

Strat, T. M. (1994), "Decision Analysis Using Belief Functions," in *Decision Making under Dempster–Shafer Uncertainties,* M. Fedrizzi, J. Kacprzyk, R. R. Yager, Eds., John Wiley & Sons, New York.

Stukey, E., and Zeckhauser, R. (1978), "Decision Analysis," in *A Primer for Policy Analysis,* W. W. Norton, New York, pp. 201–254.

Svenson, O. (1990), "Some Propositions for the Classification of Decision Situations," in *Contemporary Issues in Decision Making,* K. Borcherding, O. Larichev, and D. Messick, Eds., pp. 17–31, North Holland, Amsterdam.

Svenson, O., and Maule, A. J. (1993), *Time Pressure and Stress in Human Judgment and Decision Making,* Plenum, New York.

Swain, A. D., and Guttman, H. (1983), *Handbook for Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications,* NUREG/CR-1278, U.S. Nuclear Regulatory Commission, Washington, DC.

Tanner, W. P., and Swets, J. A. (1954), "A Decision Making Theory of Visual Detection," *Psychological Review,* Vol. 61, pp. 401–409.

Torrance, E. P. (1953), "The Behavior of Small Groups under the Stress Conditions of 'Survival,'" *American Sociological Review,* Vol. 19, pp. 751–755.

Tuckman, B. W. (1965), "Development Sequence in Small Groups," *Psychological Bulletin,* Vol. 63, pp. 289–399.

Tversky, A. (1969), "Intransitivity of Preferences," *Psychological Review,* Vol. 76, pp. 31–48.

Tversky, A. (1972), "Elimination by Aspects: A Theory of Choice," *Psychological Review,* Vol. 79, pp. 281–289.

Tversky, A., and Kahneman, D. (1973), "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Psychology,* Vol. 5, pp. 207–232.

Tversky, A., and Kahneman, D. (1974), "Judgment under Uncertainty: Heuristics and Biases," *Science,* Vol. 185, pp. 1124–1131.

Tversky, A., and Kahneman, D. (1981), "The Framing of Decisions and the Psychology of Choice," *Science,* Vol. 211, pp. 453–458.

Valenzi, E., and Andrews, I. R. (1973), "Individual Differences in the Decision Processes of Employment Interviews," *Journal of Applied Psychology,* Vol. 58, pp. 49–53.

Viscusi, W. K. (1991), *Reforming Products Liability,* Harvard University Press, Cambridge, MA.

von Neumann, J., and Morgenstern, O. (1947), *Theory of Games and Economic Behavior,* Princeton University Press, Princeton, NJ.

Wagenaar, W. A. (1992), "Risk Taking and Accident Causation," in *Risk-Taking Behavior,* J. F. Yates, Ed., John Wiley & Sons, New York, pp. 257–281.

Wallsten, T. S. (1972), "Conjoint-Measurement Framework for the Study of Probabilistic Information Processing," *Psychological Review,* Vol. 79, pp. 245–260.

Wallsten, T. S. (1976), "Using Conjoint-Measurement Models to Investigate a Theory about Probabilistic Information Processing," *Journal of Mathematical Psychology,* Vol. 14, pp. 144–185.

Wallsten, T. S. (1995), "Time Pressure and Payoff Effects on Multidimensional Probabilistic Inference," in *Time Pressure and Stress in Human Judgment,* O. Svenson and J. Maule, Eds., Plenum Press, New York.

Wallsten, T. S., Zwick, R., Kemp, S., and Budescu, D. V. (1993), "Preferences and Reasons for Communicating Probabilistic Information in Verbal and Numerical Terms," *Bulletin of the Psychonomic Society,* Vol. 31, pp. 135–138.

Weber, E. (1994), "From Subjective Probabilities to Decision Weights: The Effect of Asymmetric Loss Functions on the Evaluation of Uncertain Outcomes and Events," *Psychological Bulletin,* Vol. 115, pp. 228–242.

Weber, E., Anderson, C. J., and Birnbaum, M. H. (1992), "A Theory of Perceived Risk and Attractiveness," *Organizational Behavior and Human Decision Processes,* Vol. 52, pp. 492–523.

Wegner, D. (1987), "Transactive Memory: A Contemporary Analysis of Group Mind," in *Theories of Group Behavior,* B. Mullen and G. R. Goethals, Eds., Springer-Verlag, New York, pp. 185–208.

Weinstein, N. D. (1979), "Seeking Reassuring or Threatening Information about Environmental Cancer," *Journal of Behavioral Medicine,* Vol. 2, pp. 125–139.

Weinstein, N. D. (1980), "Unrealistic Optimism about Future Life Events," *Journal of Personality and Social Psychology,* Vol. 39, pp. 806–820.

Weinstein, N. D. (1987), "Unrealistic Optimism about Illness Susceptibility: Conclusions from a Community-Wide Sample," *Journal of Behavioral Medicine,* Vol. 10, pp. 481–500.

Weinstein, N. D., and Klein, W. M. (1995), "Resistance of Personal Risk Perceptions to Debiasing Interventions," *Health Psychology,* Vol. 14, pp. 132–140.

Welch, D. D. (1994), *Conflicting Agendas: Personal Morality in Institutional Settings,* Pilgrim Press, Cleveland.

Welford, A. T. (1976), *Skilled Performance,* Scott, Foresman, Glenview, IL.

Wickens, C. D. (1992), *Engineering Psychology and Human Performance,* HarperCollins, New York.

Winkler, R. L., and Murphy, A. H. (1973), "Experiments in the Laboratory and the Real World," *Organizational Behavior and Human Performance,* Vol. 10, pp. 252–270.

Winterfeldt, D. V., and Edwards, W. (1986), *Decision Analysis and Behavioral Research,* Cambridge University Press, Cambridge.

Yates, J. F., Ed. (1992), *Risk-Taking Behavior,* John Wiley & Sons, New York.

Zander, A. (1994), *Making Groups Effective,* 2nd Ed., Jossey-Bass, San Francisco.

Zimmer, A. (1983), "Verbal versus Numerical Processing of Subjective Probabilities," in *Decision Making under Uncertainty,* R. W. Scholtz, Ed., North Holland, Amsterdam.

Zey, M., Ed. (1992), *Decision Making: Alternatives to Rational Choice Models,* Sage, London.