

Beyond Expected Utility: Rethinking Behavioral Decision Research

Deborah Frisch and Robert T. Clemen

Much research in psychology has evaluated the quality of people's decisions by comparisons with subjective expected utility (SEU) theory. This article suggests that typical arguments made for the status of utility theory as normative do not justify its use by psychologists as a standard by which to evaluate decision quality. It is argued that to evaluate decision quality, researchers need to identify those decision processes that tend to lead to desirable outcomes. It is contended that a good decision-making process must be concerned with how (and whether) decision makers evaluate potential consequences of decisions, the extent to which they accurately identify all relevant consequences, and the way in which they make final choices. Research that bears on these issues is reviewed.

A central goal of behavioral decision research is to evaluate the quality of people's decisions. The most widely used method of making such an evaluation of quality is to compare people's decisions with a normative model of rational decision making. Naturally, the choice of a specific normative model against which to compare human behavior is crucial in defining the nature of the behavioral research that is done. Few such models are available to behavioral decision researchers, the most prominent being expected utility theory and its close cousin, subjective expected utility (SEU).

Initially developed by von Neumann and Morgenstern (1947), expected utility theory begins with a set of axioms relating to an individual's preferences among gambles. The essence of utility theory is a mathematical proof that shows that if a person's preferences conform to the axioms, then two important consequences follow. First, one can infer the person's values (known as a utility function) from observing her or his choices. Second, this person's choices can be described as if she or he were following a decision rule of maximizing expected utility (in which "expected" is used in the usual probability-theory sense to indicate a probability-weighted average). The axioms that underlie utility theory include such notions as connectivity (the ability to compare gambles in terms of preference), transitivity of preferences (if A is preferred to B , which in turn is preferred to C , then A must be preferred to C), and independence (roughly, if A is preferred to B , then a p chance at A is preferred to a p chance at B).

In their book, von Neumann and Morganstern (1947) assumed that the probabilities of uncertain events were given. Savage (1954) developed SEU by modifying and extending expected utility theory to circumstances in which probabilities

are not given. SEU arrives at the same maximum-expected-utility decision rule, except that the probabilities are the decision maker's personal or subjective probabilities for uncertain outcomes. For a more complete description of both expected utility and SEU, see von Winterfeldt and Edwards (1986).

It is useful to distinguish three different interpretations of the axioms of utility theory. These interpretations correspond to different roles the theory has played in research on decision making. First, the axioms can be interpreted as guidelines about the choices a fully rational person should make. This is the normative interpretation, which has resulted in elaboration of the SEU model (Fishburn, 1988; Keeney & Raiffa, 1976) and the development of generalized utility theories (Machina, 1982). This leads to the second interpretation, whereby the axioms can be used as tools for measuring or assessing people's beliefs and values (von Winterfeldt & Edwards, 1986). That is, if one assumes the validity of the axioms, then one can construct probability functions (which represent a person's beliefs) and utility functions (which represent values). This interpretation provides the basis for decision-analysis applications of SEU.

Finally, the axioms can be interpreted as hypotheses about the pattern of choices people actually make (descriptive interpretation). This interpretation leads to empirical research testing the validity of the axioms, which involves constructing sets of choices such that the axioms imply a certain pattern of choices. Typically, this is of the form "If you choose Option A in Situation 1, then the axioms imply that you should choose the corresponding Option A' in Situation 2"). This has been the prevailing paradigm for much of behavioral decision research in the past.

In the past four decades of behavioral decision research, SEU theory (or the slightly less general expected utility theory¹) has typically been the model of rational decision making against which human decision behavior has been compared. A great deal of evidence has been compiled demonstrating that people's choices deviate systematically from the model, whether the issue has been subjective judgment of uncertainty or choices

Deborah Frisch, Department of Psychology, University of Oregon;
Robert T. Clemen, College of Business Administration, University of Oregon.

We would like to thank Jon Baron, Ward Edwards, Robin Gregory, Daniel Kahneman, John Orbell, and an anonymous reviewer for very helpful comments on drafts of this article.

Correspondence concerning this article should be addressed to Deborah Frisch, Department of Psychology, University of Oregon, Eugene, Oregon 97403.

¹ In this article, we are concerned specifically with subjective expected utility. For convenience, however, we occasionally use the term *utility theory* as a synonym for *subjective expected utility*.

among alternatives. The predominant interpretation of these findings is that people fail to conform to the normative model as a result of various cognitive limitations (for reviews, see Baron, 1988; Dawes, 1988).

The use of utility theory as a standard for evaluating the quality of decisions has generated several productive lines of research. First, the fact that choices systematically deviate from utility theory has led researchers to develop alternative models that provide a more accurate description of people's choices (e.g., Bell, 1982; Fishburn, 1988; Kahneman & Tversky, 1979; Loomes & Sugden, 1982; Lopes, 1987; Machina, 1982; Tversky & Kahneman, 1992). There has also been a great deal of research attempting to describe the cognitive processes involved in decision making that give rise to these violations (Payne, Bettman, & Johnson, 1992). Finally, some researchers have argued that people can and should be trained to conform to utility theory. On the basis of the normative model (what one ought to do) and the descriptive violations of it (what one actually does), prescriptive techniques can be developed to bring the "actual closer to the ideal" (Kleinmuntz, 1991). This is the view explicitly endorsed by many psychologists and practicing decision analysts (Baron, 1988; Bell, Raiffa, & Tversky, 1988; Kleinmuntz, 1991).

Thus, the use of utility theory as a standard for evaluating the quality of decisions has generated a great deal of empirical research examining how people make decisions and some practical advice for improving decisions. The purpose of this article is to examine critically the extent to which utility theory provides an adequate standard by which to evaluate decision quality. Our thesis is that utility theory, as it is typically used to provide guidance for behavioral decision research, is not adequate either as a description of how decisions are made or as a standard of good decision-making behavior.

We are not the first authors to critique the use of utility theory as a standard of decision making (e.g., see Hastie, 1991; Lopes, 1981). However, our approach differs from other critiques of utility theory in two important ways. First, whereas most other critics have questioned particular assumptions of utility theory, we argue that utility theory is not the right type of model for psychologists interested either in describing people's decisions or in prescribing how people might make better decisions. Second, some critics have suggested that psychological research on decision making should describe people's behavior and should not attempt to evaluate the quality of decisions. In contrast, we believe that a useful research strategy is to compare people's behavior with a standard. An agreed-on standard of decision making leads to a coherent research agenda and provides a way to organize empirical findings (Hastie, 1991). Moreover, this research strategy allows for the possibility that one's empirical research can lead to practical applications. By discovering ways that people go astray in decision making, one not only gains insight into psychological processes but also has a chance of improving decision making. For this reason, an important goal of this article is to develop a new framework for evaluating decision quality.

The remainder of the article is organized into three sections. First, we review the arguments that are typically used to justify the use of utility theory as a standard of decision making. We argue that utility theory is not an adequate model for psycholo-

gists, either for providing a standard for evaluating the quality of decisions or for describing how people make decisions. Second, we propose an alternative framework for defining good decision making. We show how this framework generates a variety of new empirical questions, providing a new and insightful way to organize some recent behavioral research in decision making. Finally, we discuss the relationship between our framework and other approaches to decision making.

Descriptive and Normative Functions of SEU

As described earlier, SEU theory is meant to provide a normative guide for an individual decision maker. However, Savage (1954) originally proposed this theory as both descriptively accurate and normatively compelling. In what sense could utility theory be a descriptive model of decision making? If the axioms were empirically true, then one could describe a person's behavior as if the rule of maximizing SEU was being followed. This sense in which utility theory could be descriptively accurate is very narrow. At best, the model would provide a description of the patterns of choices people make. It would not describe the psychological processes involved in decision making. It would not, for example, distinguish situations in which one consciously chose an option from situations in which one acted out of habit or impulsiveness. However, we suspect that most psychologists believe that a descriptive model should describe the psychological processes involved in decision making. If this is the case, then the sense in which utility theory might be descriptive is quite different from what most psychologists want from a descriptive theory of decision making.

Of course, even though Savage (1954) did not intend for utility theory to describe the psychological processes involved in decision making, it is possible (although rarely done) to interpret the theory as a process model. Specifically, one might interpret SEU as postulating that people assess probabilities and utilities, compute expected utility, and choose the option with the highest expected utility. If one wanted to test the process interpretation of utility theory, then presumably one would want to determine whether subjects behaved in the prescribed manner. However, this is not the type of research that is typically conducted by behavioral decision researchers. In general, tests of utility theory examine whether people's choices conform to the axioms of the theory. Subjects are presented with hypothetical decision scenarios and are asked to make a choice. These scenarios are usually constructed such that the axioms of utility theory require a certain pattern of choices (Frisch, 1993). If the pattern of choices observed is inconsistent with utility theory, then the study is considered to be evidence against the empirical validity of the theory. It is clear, however, that such evidence primarily addresses SEU's adequacy as an as-if model and says little about its adequacy as a process model.

There has been some research examining the extent to which people make decisions consistent with the process interpretation of SEU. Furby and Beyth-Marom (1991) described a model of the decision-making process based on SEU and organized research on adolescent risk taking within that framework. They concluded that existing studies on adolescent decision making do not provide conclusive evidence about whether adolescents' decisions are consistent with the model. Janz and Becker (1984)

reviewed a large body of research on the health belief model (HBM), a model of decisions about compliance with various preventive-health behaviors (e.g., regular physical activity) and sick-role behaviors (e.g., taking one's medication). Janz and Becker (1984) concluded that the HBM, which is conceptually very similar to SEU, provides a reasonably good predictive model of the extent to which people will engage in health-related behaviors.

Even if utility theory is interpreted as being descriptive of the psychological processes involved in decision making, many researchers (e.g., Baron, 1988; Hastie, 1991) have noted that it describes only a very small part of the process. Specifically, utility theory describes how one makes a decision once it is structured but does not describe how one generates options, determines which consequences to consider, or identifies the relevant risks. Decision analysts refer to this as the "up-front" or "structuring" part of the decision-making process, and in practice it is often the most important and difficult part of a decision (e.g., see Clemen, 1991; Keeney, 1992; Polister, 1991; or von Winterfeldt & Edwards, 1986).

Thus, when researchers say that utility theory fails descriptively, they mean that there are conditions under which the axioms are violated. The fact that people violate the axioms of utility theory obscures the more important point that, in principle, utility theory is not the type of descriptive theory psychologists need. Indeed, many researchers have ignored utility theory and instead have attempted to describe the processes involved in decision making (see Payne, Bettman, & Johnson, 1992, for a review).

Similar problems arise when one examines the sense in which utility theory could be a standard for evaluating the quality of people's decisions. Two types of arguments are typically used to justify utility theory as a normative model. The *logical argument* starts with a set of seemingly intuitive, uncontroversial axioms and logically derives a specific decision rule. If one accepts the axioms, then one is compelled to accept the conclusion. Savage (1954) described this quite clearly:

In what sense is this theory normative? It is intended that a reflective person who finds himself about to behave in conflict with the theory will reconsider. . . . To use the preference theory is to search for incoherence among potential decisions, of which you, the user of the theory, must then revise one or more. The theory itself does not say which way back to coherence is to be chosen, and presumably should not be expected to. (p. 308)

Thus, violations of SEU imply internal inconsistency. Other authors have shown that if one's preferences are inconsistent with SEU, then one is susceptible (in principle) to exploitation by more rational decision makers. This is known as the *coherence argument*, which consists of a demonstration that violations of the rule can lead to undesirable consequences in which the decision maker is subject to a sure loss (e.g., see Bunn, 1984, chap. 2, for examples of money pumps and Dutch books).

Thus, SEU theory is normative in the following sense: If one accepts the axioms, then, to be internally consistent, one's choices must conform to the rule of maximizing expected utility. Debate about the normative status of SEU theory has questioned whether one should, in fact, accept the axioms (Shafer, 1986; Slovic & Tversky, 1974). This debate has obscured the fact that even if one accepts the axioms, utility theory

is normative in a very narrow sense. If people conform to utility theory, their decisions are internally consistent. It is not clear, however, that internal consistency (i.e., consistency with the axioms of utility theory) is either a necessary or a sufficient criterion for good decision making.

Several researchers have questioned whether conformity to the axioms of utility theory is a necessary component of good decision making. For example, some researchers, such as Lopes (1981), have objected to the particular rule of maximizing expected utility. Lopes argued that a normative model should not require maximization but should acknowledge the fact that people have other legitimate goals in decisions (e.g., security). Similarly, Loomes and Sugden (1982) claimed that the axioms of utility theory "constitute an excessively restrictive definition of rational behavior" (p. 823).

There is another objection to the claim that conformity to the axioms is a necessary component of good decisions. Some researchers have argued that the specific assumptions of utility theory are not normatively compelling (Allais, 1953; Ellsberg, 1961; Frisch & Jones, 1993; Machina, 1989; Slovic & Tversky, 1974). For example, consider the "framing effect" phenomenon, which refers to the finding that people's choices sometimes vary as a function of how a situation is described or framed (Kahneman & Tversky, 1984). The framing effect is one of the most serious violations of utility theory that has been demonstrated. However, there is some empirical evidence demonstrating that frame can influence people's experience of the outcome of decisions (Levin & Gaeth, 1988). For example, Levin and Gaeth (1988) demonstrated that subjects who ate ground beef described as "75% lean" reported a more favorable experience than subjects who ate the same ground beef described as "25% fat." If framing influences one's experience of the consequences of a decision, then it seems reasonable for one's decision to vary as a function of frame (see Bell et al., 1988; Frisch & Jones, 1993; Kahneman & Tversky, 1984). Thus, conformity to the axioms of utility theory might not be a necessary component of good decision making.²

A more serious problem is that utility theory does not provide a sufficient standard of good decision making. Imagine that empirical research had demonstrated that people's choices do conform to the axioms of utility theory. Would researchers conclude that people are making the best possible decisions for themselves? Would such a finding imply that the outcomes of people's decisions are desirable? We think not. Even if people conform to utility theory, they may be doing a poor job of specifying the uncertainties they face or understanding the consequences of their actions, for example. In fact, there may be many ways in which an expected utility maximizer might im-

² This example need not be interpreted as evidence against the normative status of the axioms of utility theory. One might argue that "eating beef described as 75% lean" is a different event than "eating beef described as 25% fat." More generally, any apparent violation of the axioms of utility theory can be explained by arguing that the consequences were actually different in the two cases. However, if one makes this argument, the theory quickly becomes one that requires situation-specific definitions of preferences. Thus, the theory loses much of its explanatory and predictive value (see Machina, 1989, for a detailed discussion of this issue).

prove his or her overall decision process. Kahneman and Snell (1992) made a similar point by distinguishing between “decision utility,” which is inferred from a person’s choices, and “experience utility,” which refers to the person’s actual experience of the consequences. In Kahneman and Snell’s (1992) terminology, we are claiming that a person who maximizes decision utility is not necessarily maximizing experience utility.

From our perspective, a standard of decision making provides guidelines about how one should make decisions so as to balance the desirability of outcomes against the chance of obtaining them. Again, one might interpret SEU as stating that “the process of decision making that balances desirability and likelihood of outcomes involves assessing probabilities and utilities and then combining them to make a choice.” In fact, we do believe that this interpretation provides a starting point for a model of a good decision-making process. The reasons, however, have nothing to do with internal consistency or the compellingness of the axioms. Rather, we believe that SEU is justified as a starting point because it does provide a way to achieve a balance once desirability and likelihood of consequences have been determined.

In summary, there has been much debate about whether the axioms of utility theory are descriptively valid and whether they are normatively compelling. Our point is that even if the axioms are accurate and even if researchers agree that they are normatively compelling, many unanswered questions remain about how people make decisions and whether they make good decisions. Whether one is interested in describing the way people actually make decisions or in prescribing how people can make better decisions, one needs a process model of decision making. Clearly, utility theory is not (nor was it intended to be) descriptive of the decision-making process and does not provide an adequate standard for this process.

An Alternative Framework

We have argued that utility theory is not an adequate standard for researchers interested in evaluating the quality of decisions. In contrast, we have suggested that an adequate standard would define the components of a good decision-making process. How would one go about developing and justifying a model of a good decision-making process? In this section, we propose the foundations of such a model and describe the methods one would use to justify it. Our approach is based on the idea that decision outcomes should be used to define a model of a good decision-making process.

Of course, as decision researchers often note, the desirability of the outcome of a particular decision does not provide conclusive evidence about the quality of that decision. Good decisions can lead to bad outcomes as a result of uncertainty in the environment. In the case of uncertainty, we postulate that good decision-making processes tend (on average) to lead to more desirable outcomes than do poor decision-making processes. Thus, one could empirically derive a model of a good decision-making process by comparing many decisions with positive and negative outcomes and identifying systematic differences.

In this section, we describe the foundations of such a model and the empirical questions that arise from it. In our view, there are two classes of empirical questions that should be addressed

about each stage of the decision-making process. The first question is whether, in fact, a particular component is associated with positive outcomes. The second question examines the extent to which people actually engage in each process. We discuss how some recent research in decision making provides partial answers to these questions. We also highlight several important empirical questions that remain relatively unexplored.

We begin with three basic features that, we propose, must be a part of good decision making. 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 trade-offs of some form (*compensatory decision rule*). We argue that these three components are important in the sense that they help the decision maker achieve better outcomes.

Consequentialism

The first component of a good decision is that one should choose a course of action on the basis of the expected consequences of different actions.³ That is, consequentialist decision strategies are preferable to nonconsequentialist ones. This claim is based on the notion that, to achieve a balance between the desirability and likelihood of outcomes, one’s decision process should focus explicitly on the consequences of different actions.

It is useful to distinguish two related empirical questions one might ask about consequentialism. The first question is whether a decision is actually based on a consideration of the desirability and likelihood of outcomes. This is a question about the decision-making process. As Dawes (1988) noted, decisions can be made on the basis of a variety of nonconsequentialist arguments, including habit, tradition, and imitation. The second question is whether a decision is consistent with one’s assessments of the desirability of outcomes. Common sense suggests that there are situations in which a person’s choices do not reflect these assessments. For example, imagine that a person has a \$1,000 balance on his or her VISA card and pays 17% annual interest. Imagine that this person also has \$2,000 in a savings account earning 4% interest. From the perspective of utility theory, one would conclude that the utility to this person from having money in a savings account (e.g., feelings of security) outweighed the cost of paying the high interest on the VISA bill. Our claim is that this is possible but is not necessary. That is, it is an empirical question whether and when people’s choices reflect their assessments of the desirability of different outcomes. In this situation, it is possible, for example, that the person developed the habit of having a debt on the VISA bill when he or she had little money in savings or when the interest rate on savings was much higher. Thus, the current decision might be

³ The term *consequentialism* has slightly different meanings in different contexts. Here, we use the term informally and without a precise definition to refer to a decision strategy in which the choice of an option is based on one’s evaluation of its expected consequences. We mean to contrast this approach with decision strategies in which the choice of an option is based on other kinds of (possibly nonevaluative) behavior such as imitation of peers or habit.

based on (a no longer justified) habit and not on a careful consideration of the consequences.

Neither of these issues comes up in discussions of typical normative models such as SEU theory. Because SEU is an as-if model, it is irrelevant (in fact meaningless) to talk about whether people actually base decisions on future consequences (Becker, 1976). The typical approach also assumes that people's choices reflect their assessments of the likelihood and desirability of outcomes. Because preferences (i.e., assessments of the desirability of outcomes) are inferred from observed choices, it does not make sense to say that choices are inconsistent with preferences.

From our perspective, there are two basic empirical questions about consequentialism. First, to what extent and under what conditions do people make decisions on the basis of an assessment of the likelihood and desirability of different possible outcomes? Second, are decisions always consistent with the decision maker's assessments of the desirability of outcomes?

Recent research by Tversky and Shafir (1992) suggests that people sometimes base decisions on reasons that are nonconsequentialist. Subjects were asked whether they would like to spend a vacation in Hawaii under a variety of conditions. Some subjects who stated that they would want to spend a vacation in Hawaii regardless of the outcome of an examination (pass or fail) also stated that they would rather postpone the decision if the outcome of the exam was unknown. Consequentialist reasoning would suggest that if a person prefers to go to Hawaii whether she or he fails or passes, then she or he should prefer to go when the outcome is unknown. Baron (1994) also has described some empirical findings in which people's choices are nonconsequentialist. In particular, he described findings in which people's choices are inconsistent with their assessments of the desirability of different outcomes.

Thus, there is evidence demonstrating that people sometimes make decisions based on nonconsequentialist reasoning. There are a variety of relatively unexplored empirical questions related to this issue. The most obvious and important empirical question is "Under what conditions do people make decisions on the basis of careful consequentialist thought, and when do they make decisions using other types of decision rules?" Are there individual differences in the extent to which consequentialist decision rules are used? It is plausible to expect so. For example, formal education and training might have an effect on the use of consequentialist decision rules. Larrick, Morgan, and Nisbett (1990) found that training in cost-benefit thinking affected people's reasoning about everyday decisions. Does the use of consequentialist decision rules vary across content domains? A recent study of real-life decision making (Frisch, Jones, & O'Brien, 1993) found that people were more likely to use consequentialist strategies for professional decisions (e.g., career or education) than for personal decisions (e.g., intimate relationships).

There are also a variety of empirical questions that can examine the validity of our claim that consequentialist decision strategies are preferable to nonconsequentialist ones. Are people more satisfied with decisions made using such strategies? Are good outcomes associated with consequentialist strategies and bad outcomes associated with nonconsequentialist ones? Jones and Frisch (1993) have found evidence for an affirmative

answer. Subjects were asked to describe either good or bad outcomes that had occurred in their lives. They were also asked to describe the thinking that led to the outcomes. Consequentialist thinking was reported more often by those subjects describing good outcomes than by those describing bad outcomes.

Thus, there are many relatively unexplored empirical questions about the conditions under which people make decisions based on consequentialist thinking and the extent to which this type of thinking is associated with good outcomes. Although these are basic questions about decision making, they do not make sense from the perspective of utility theory. Utility theory assumes that decisions reflect a person's beliefs about the likelihood and desirability of possible consequences. Empirical tests of utility theory (in any form) typically do not examine this assumption but are designed to determine whether a decision maker's choices are consistent, given the specific form of consequentialism implicitly assumed in the theory.

Thorough Structuring

To make good decisions, it is not sufficient to think about the consequences of one's actions; one must do so accurately and thoroughly. The second component of a good decision is that one should try to accurately anticipate the different possible consequences of one's actions. To anticipate the consequences of one's actions accurately, one must predict accurately both what consequences will (or might) occur and how one will experience the different possible consequences. We use the term *structuring* to refer to the process of identifying the possible actions one might take, identifying the different consequences of one's actions, and assessing the likelihood and desirability of those consequences. Thus, structuring a decision can be decomposed into (a) option generation, (b) belief structuring (determining what will or might happen as a result of different options), and (c) value structuring (determining the desirability of the different possible consequences of different options).

How should one structure a decision? Much empirical research is needed to clarify the aspects of structuring that are most relevant to achieving good outcomes. We describe here some basic issues involved with each of the three components.

Guidelines for generating options have been provided by Keller and Ho (1988), Keeney (1988, 1992), and Gregory and Keeney (1992). An important point is that one should explicitly consider and evaluate the consequences of more than one option. An implication of this claim is that when a decision involves a choice between the status quo and some alternative, one should view the status quo as an option. That is, it is an illusion to think that one can "do nothing" or "not make a decision." Maintaining the status quo should be viewed as an option that has certain consequences, and that option should be examined as carefully as any other option.

Belief structuring refers to anticipating the potential consequences of one's actions and is necessary in decisions made under both certainty and uncertainty. For example, imagine that a person goes to a Chinese restaurant and orders the curry tofu. Unfortunately, the curry tofu contains a large quantity of chili peppers, which the person dislikes. We could characterize the decision as a failure of belief structuring because the person failed to accurately anticipate the consequences of the decision.

In contrast, value structuring refers to determining the desirability of consequences. If the individual in the previous example had known that the curry tofu contained chili peppers but had not realized that she or he disliked them, this would have been a failure of value structuring. Value structuring requires considerable thought and introspection about one's goals and objectives. In some cases, it may be very difficult to anticipate and evaluate consequences. Keeney (1992) provided some prescriptive advice in this regard.

There are many empirical questions about how people structure decisions. How many options do people typically generate in making a decision? When decisions involve a choice between the status quo and a change, do people view the status quo as an option and evaluate the consequences of maintaining the status quo? Jones and Frisch (1993) found evidence that people think differently about decisions that are framed as choices (e.g., Should I save my money or should I buy a bicycle?) as opposed to those that are framed as opportunities (e.g., Should I buy a bicycle?).

Are certain classes of consequences systematically ignored (e.g., long-term consequences)? Are certain classes of consequences overweighted (e.g., short-term emotional consequences)? To what extent are surprising or unexpected consequences due to a failure to predict accurately what will happen versus a failure to predict accurately how one will experience the consequence? Kahneman and his colleagues (Kahneman & Snell, 1992; Varey & Kahneman, 1992) have recently begun to examine the extent to which people accurately anticipate the consequences of their actions. This research falls in the category we call value structuring. They have identified several potential areas in which people might make inaccurate predictions about the desirability of outcomes. For example, subjects were not very accurate in predicting the effect of repeated exposure to a stimulus (e.g., yogurt) on the future desirability of that stimulus.

Finally, how do people decide how much structuring to do and when to stop? Do people tend to structure certain kinds of decisions more carefully than others? Certainly some decisions deserve more careful treatment in this respect than others. What sort of advice can be given to help a decision maker manage the structuring process in this regard? Phillips (1984), for example, defined the notion of a requisite decision model, one that contains everything that is essential to solving the problem but ignores nonessential issues. Sensitivity analysis often can be used to determine what is essential and what is not (Clemen, 1991).

Although we have described some features that we think should ideally characterize the structuring process, there is a need for empirical research examining the aspects of structuring that are associated with good versus bad outcomes. What is the relationship between the number of options considered and the quality of the outcome? To what extent can bad outcomes be attributed to failures of belief or value structuring? Answering such questions will help to define aspects of good decision making.

Compensatory Decision Rule

The third component of a good decision is that when one's actions may have several different possible consequences, one

should attempt to make trade-offs. A decision should reflect the desirability and likelihood of all of the different consequences of different possible actions. We are not concerned here with the particular rule used to make trade-offs. Rather, we wish to stress the more general point that decisions based on compensatory rules are preferable to those based on noncompensatory rules.

Why is a compensatory decision rule preferable to a noncompensatory one? One argument is basically the same as that for consequentialism and complete structuring: One is more likely to achieve desirable outcomes if one's decisions reflect all of the possible consequences of one's actions. Even authors who reject the particular rule of maximizing SEU endorse some type of compensatory decision rule (e.g., Lopes, 1990).

Empirical questions in regard to the compensatory decision rule include the following: Under what conditions are compensatory decision rules used? When do people make trade-offs? Research on this topic suggests that the relative weight given to different components of a decision varies as a function of how the choice is presented (Payne et al., 1992; Tversky, Sattath, & Slovic, 1988). Moreover, there is evidence that people try to avoid making trade-offs if possible. Montgomery (1983) and Svenson (1992) have argued that people attempt to reframe decisions so that they do not need to make trade-offs. What kinds of trade-offs do people perceive as difficult? Goldstein and Beattie (1991) suggested that trade-offs are difficult when one must compare attributes that one perceives are similar in relative importance.

Again, there are interesting empirical questions examining whether the use of a compensatory decision rule is a necessary component of a good decision. Are good outcomes more likely when people explicitly made trade-offs in the decision-making process? Do compensatory rules tend to lead to more desirable outcomes than noncompensatory ones? Payne, Bettman, and Johnson (1988) provided an analysis of the conditions under which compensatory strategies are preferred to noncompensatory ones. They found that people's selection of a decision strategy was sensitive to the costs and benefits of different strategies.

Relation to Previous Research

The framework we have described suggests a way of conceptualizing a good decision that is different from the traditional, utility theory approach. Moreover, our approach provides a way to organize a variety of empirical findings in decision making and suggests some unexamined empirical questions. This approach has much in common with that of Janis and Mann (1977) and is also consistent with much recent work in decision analysis. In this section, we discuss these connections.

Relation to Janis and Mann

Janis and Mann (1977) described features of "vigilant" processing, which they proposed as an ideal decision-making process. Vigilant processing includes many of the features of the process we have described, particularly those involved in thorough structuring. Like Janis and Mann, we believe that a criterion by which to evaluate the quality of decision making should consist of a description of the processes by which decisions should be made. This is in contrast to the traditional approach,

which uses an axiomatic model as a standard. Also, Janis and Mann's approach is based on the premise that the ultimate justification for a standard of good decision making is empirical. That is, they proposed that the process model described will tend to lead to desirable outcomes, but they viewed this as an empirical question. Similarly, in our view the ultimate justification of a standard of decision making is empirical. This is in contrast to the SEU approach, which is based on the premise that the justification of a normative model is logical or mathematical.

We have argued that the justification of a particular set of processes as a standard must come from a demonstration that those processes tend to be associated with desirable outcomes in the real world. Like Janis and Mann, we recognize that there are serious difficulties with attempting to relate the decision-making process to the quality of outcomes. It is very difficult to evaluate the quality of a decision's outcome. Most decisions have desirable and undesirable consequences. Furthermore, there are methodological problems involved in attempting to determine the processes involved in real-life decisions, particularly with respect to the validity of subjects' self-reports about how they approached the decision.

In spite of these difficulties, some research suggests that it is feasible to examine the relationship between the decision-making process and the desirability of outcomes. Herek, Janis, and Huth (1987) attempted to investigate empirically whether desirable outcomes were associated with the use of the processes outlined in Janis and Mann (1977) by examining presidential decision making in international crises. Experts rated both the quality of the decision-making process and the quality of the outcome and found a strong relationship between the two measures. Larrick, Nisbett, and Morgan (1993) found a relationship between the use of certain normative decision principles (e.g., the sunk cost principle) and the quality of outcomes. They used objective measures such as salary (for faculty) and grade point average (for students) to define the quality of outcomes. The work by Jones and Frisch (1993) described earlier used subjects' self-reports about the decision-making process and the quality of the outcomes of decisions.

Thus, there are a variety of ways researchers can operationally define a good outcome. Independent coders might be used to assess the quality of the outcomes. Objective measures such as grade point average or salary might be used as part of the definition of good outcomes. Subjects' subjective assessments of the desirability of outcomes represent another way of defining a good outcome. Our intention is not to defend any particular method but to stress the importance of examining the relationship between the decision-making process and the quality of outcomes.

Relation to SEU and Decision Analysis

As we stated earlier, the view we have presented is not incompatible with SEU theory. Rather, our view emphasizes the decision-making process, whereas SEU theory focuses on the consistency of choices. In this section, we describe the relationship and the crucial differences between the components of the framework we have described and SEU theory.

The idea of consequentialism is implicit in SEU theory. SEU

theory assumes that people base decisions on the expected consequences of different actions. Our approach makes the consequentialist assumption explicit and leads to two basic empirical questions that mainstream behavioral decision theory has largely ignored. First, under what conditions do people base decisions on consequentialist reasons as opposed to other types of reasons? Second, are consequentialist decision strategies more strongly associated with good outcomes than nonconsequentialist strategies?

The idea of thorough structuring deals with an aspect of the decision-making process that is not directly addressed in SEU theory, which prescribes how one should choose among options given one's beliefs and values. It does not prescribe which options one should consider or the consequences that should matter. SEU theory indirectly specifies that certain factors should not influence the desirability of consequences because the axioms place constraints on the pattern of preferences that are coherent. For example, the axioms of SEU imply that the desirability of a consequence should be independent of the description (frame) of the decision situation. Although SEU theory does not specify what should influence decisions (Dawes, 1988), the underlying axioms do imply that the structuring process can be divided into values, on one hand, and beliefs about uncertainty, on the other. Our framework focuses on the question of how people structure decisions and which aspects of structuring are most strongly associated with the quality of the outcome.

Our claim that decisions should be based on a compensatory decision rule is consistent with but weaker than SEU. SEU prescribes a particular compensatory decision rule—mathematical expectation of utility—whereas other axiom-based theories lead to alternative specific decision rules. Our approach focuses on the question of whether any compensatory decision rule is used and whether the rule used is related to the quality of the outcome.

Although nothing in our approach directly contradicts SEU theory and the two perspectives overlap substantially, there are also several important differences. Utility theory was developed at a time when psychological research was grounded in behaviorism. Indeed, much of the appeal of SEU to psychologists stemmed from the fact that it allowed researchers to develop rich theory based solely on observed choices. Although psychologists no longer explicitly endorse behaviorism, behaviorist assumptions still are implicit in typical empirical research on decision making. In particular, although the cognitive revolution in psychology (Hintzman, 1993) has had a tremendous impact on descriptive decision research, it has had no effect on the overall structure of normative theories. It is true that empirical findings have led many researchers to question the validity of the axioms and to develop generalized utility theories on the basis of weaker axiom systems. However, generalized utility theories remain axiomatic in nature; our thesis is that behavioral research can benefit from a normative standard that is based more on realistic cognitive processes than on axiomatic systems.

Traditionally, normative models of decision making are justified by means of logical or mathematical arguments. In contrast, our central thesis is that the justification of a standard of good decision making must rely on empirical findings that

identify the aspects of the decision-making process that tend to lead to desirable outcomes. We have described several components of the decision-making process that are plausible candidates for a standard of good decision making, and it is ultimately an empirical question whether these components actually are associated with desirable outcomes. Thus, our way of defining and evaluating a standard of decision making is very different from the typical approach.

Our approach also focuses attention on different types of empirical questions. Instead of asking whether people's choices are consistent with a formal model, our approach focuses on the processes involved in decision making. Some of the questions raised by our approach have already been addressed by researchers in the field. Thus, although not all of the issues we have raised are entirely new, we believe that it is useful to provide a comprehensive framework for generating and synthesizing empirical research examining the decision-making process.

Finally, from our perspective, although SEU theory is not an appropriate standard for behavioral decision researchers to use in evaluating the quality of decisions, it can be viewed as a useful tool for improving decision making in certain situations. SEU appropriately focuses a decision maker's attention on consequences, explicitly identifies the important distinction between beliefs and values, provides guidance (within the constraints of the model) on probability and utility assessment, and provides a compensatory decision rule. Our objective in using SEU would be to help people make decisions that balance the desirability and likelihood of consequences instead of the narrower objective of ensuring that the decision maker conforms to a set of behavioral axioms. In fact, this prescriptive orientation clearly opens the door for much more in the way of decision-making advice than can be gleaned from the axioms of expected utility or other similar axiom-based models. It is no surprise that decision analysts rely heavily on techniques that fall to some extent beyond the limits of SEU—notably structuring and sensitivity analysis—to help their clients make improved decisions. Moreover, our approach indicates the importance of empirical research to test the validity (i.e., the tendency to lead to more preferred outcomes) of such decision-analysis techniques.

Conclusions

In reviewing and assessing the behavioral decision-making literature, we have concluded that utility theory is not the right type of model for researchers interested in either describing how people make decisions or evaluating the quality of people's decisions. Thus, in this article, we have suggested that psychological research on decision making can usefully be organized around a model that describes a good decision-making process.

We have outlined three criteria that are plausible candidates for components of a decision-making process that will tend to lead to good outcomes by balancing the desirability and likelihood of possible outcomes. We propose that this balance can be achieved if (a) decisions are based on the possible consequences of actions, (b) assessments of the likelihood and desirability of consequences are accurate, and (c) decisions are based on trade-offs among these factors. Although these criteria are quite plausible, we are proposing that empirical research will determine

whether they are appropriate components for a standard of good decision making, a claim that research to date appears to support. Still, many questions remain unexplored.

Although we have argued that SEU theory is an inappropriate standard by which to evaluate the quality of decisions, we have maintained the premise that empirical research is usefully organized around a standard of good decision making. Our approach offers the potential for richer descriptive models of decision making and for useful prescriptive advice for improving decisions.

References

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine [The behavior of rational man in risk situations: A critique of the axioms and postulates of the American School]. *Econometrica*, 21, 503–546.
- Baron, J. (1988). *Thinking and deciding*. Cambridge, England: Cambridge University Press.
- Baron, J. (1994). Non-consequentialist decisions. *Behavioral and Brain Sciences*, 17, 1–42.
- Becker, G. (1976). *The economic approach to human behavior*. Chicago: University of Chicago Press.
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30, 961–981.
- Bell, D. E., Raiffa, H., & Tversky, A. (1988). Descriptive, normative, and prescriptive interactions in decision making. In D. E. Bell, H. Raiffa, & A. Tversky (Eds.), *Decision making: Descriptive, normative and prescriptive interactions*. (pp. 9–30). Cambridge, England: Cambridge University Press.
- Bunn, D. (1984). *Applied decision analysis*. New York: McGraw-Hill.
- Clemen, R. T. (1991). *Making hard decisions: An introduction to decision analysis*. Boston: PWS-Kent.
- Dawes, R. M. (1988). *Rational choice in an uncertain world*. San Diego, CA: Harcourt Brace Jovanovich.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Fishburn, P. C. (1988). *Nonlinear preference and utility theory*. Baltimore: Johns Hopkins University Press.
- Frisch, D. (1993). Reasons for framing effects. *Organizational Behavior and Human Decision Processes*, 54, 399–429.
- Frisch, D., Jones, K. T., & O'Brien, M. (1993). *Sources of difficulty in real-life decision making*. Manuscript submitted for publication.
- Frisch, D., & Jones, S. K. (1993). Assessing the accuracy of decisions. *Theory and Psychology*, 3, 115–135.
- Furby, L., & Beyth-Marom, R. (1991). Risk taking in adolescence: A decision-making perspective. *Developmental Review*, 12, 1–44.
- Goldstein, W. M., & Beattie, J. (1991). Judgments of relative importance in decision making: The importance of interpretation and the interpretation of importance. In D. R. Brown & J. E. K. Smith (Eds.), *Frontiers in mathematical psychology* (pp. 110–137). New York: Springer-Verlag.
- Gregory, R., & Keeney, R. (1992). *Creating policy alternatives using stakeholder values*. Working paper, Decision Research, Eugene, OR.
- Hastie, R. (1991). A review from a high place: The field of judgment and decision making as revealed in its current textbooks. *Psychological Science*, 2, 135–138.
- Herek, G. M., Janis, I. L., & Huth, P. (1987). Decision making during international crises: Is quality of process related to outcome? *Journal of Conflict Resolution*, 31, 203–226.
- Hintzman, D. L. (1993). 25 years of learning and memory: Was the cognitive revolution a mistake? In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV* (pp. 359–391). Cambridge, MA: MIT Press.

- Janis, I., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment*. New York: Free Press.
- Janz, N. K., & Becker, M. H. (1984). The health belief model: A decade later. *Health Education Quarterly*, *11*, 1–47.
- Jones, S. K., & Frisch, D. (1993). *Using outcome information to evaluate decision quality*. Unpublished manuscript.
- Kahneman, D., & Snell, J. (1992). Predicting a changing taste: Do people know what they will like? *Journal of Behavioral Decision Making*, *5*, 187–200.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.
- Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, *39*, 341–350.
- Keeney, R. L. (1988). Value-focused thinking and the study of values. In D. E. Bell, H. Raiffa, & A. Tversky (Eds.), *Decision making: Descriptive, normative and prescriptive interactions* (pp. 465–494). Cambridge, England: Cambridge University Press.
- Keeney, R. L. (1992). *Value-focused thinking*. Cambridge, MA: Harvard University Press.
- Keeney, R. L., & Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value tradeoffs*. New York: Wiley.
- Keller, L. R., & Ho, J. L. (1988). Decision problem structuring: Generating options. *IEE Transactions on Systems, Man, and Cybernetics*, *18*, 715–728.
- Kleinmuntz, D. N. (1991). Decision making for professional decision makers. *Psychological Science*, *2*, 135–141.
- Larrick, R. P., Morgan, J. N., & Nisbett, R. E. (1990). Teaching the use of cost-benefit reasoning in everyday life. *Psychological Science*, *1*, 362–370.
- Larrick, R. P., Nisbett, R. E., & Morgan, J. N. (1994). Who uses the cost-benefit rules of choice? Implications for the normative status of economic theory. *Organizational Behavior and Human Decision Processes*, *56*, 331–347.
- Levin, I. P., & Gaeth, G. J. (1988). How consumers are affected by the framing of attribute information before and after consuming the product. *Journal of Consumer Research*, *15*, 374–378.
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, *92*, 805–824.
- Lopes, L. L. (1981). Decision making in the short run. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 377–385.
- Lopes, L. L. (1987). Between hope and fear: The psychology of risk. *Advances in Experimental Social Psychology*, *20*, 255–295.
- Lopes, L. L. (1990). Re-modeling risk aversion: A comparison of Bernoullian and rank dependent value approaches. In G. M. von Furstenberg (Ed.), *Acting under uncertainty: Multidisciplinary conceptions* (pp. 267–299). Norwell, MA: Kluwer Academic.
- Machina, M. J. (1982). 'Expected utility' analysis without the independence axiom. *Econometrica*, *50*, 277–323.
- Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, *17*, 1622–1668.
- Montgomery, H. (1983). Decision rules and the search for dominance structure: Towards a process model of decision making. In P. C. Humphreys, O. Svenson, & A. Vari (Eds.), *Analysing and aiding decision processes* (pp. 343–369). Amsterdam: North-Holland.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 534–552.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavior decision research: Theory and application. *Annual Review of Psychology*, *43*, 87–131.
- Phillips, L. (1984). A theory of requisite decision models. *Acta Psychologica*, *56*, 29–48.
- Politzer, P. E. (1991). Structuring decision analyses: Statistical and psychological considerations [Editorial]. *Journal of Behavioral Decision Making*, *4*, 79–82.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Shafer, G. (1986). Savage revisited. *Statistical Science*, *1*, 463–485.
- Slovic, P., & Tversky, A. (1974). Who accepts Savage's axioms? *Behavioral Science*, *19*, 368–373.
- Svenson, O. (1992). Differentiation and consolidation theory of human decision making: A frame of reference for the study of pre- and post-decision processes. *Acta Psychologica*, *80*, 143–168.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representations of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.
- Tversky, A., Sattath, S., & Slovic, P. (1988). Contingent weighting in judgment and choice. *Psychological Review*, *95*, 371–384.
- Tversky, A., & Shafir, E. (1992). The disjunction effect in choice under uncertainty. *Psychological Science*, *3*, 305–309.
- Varey, C., & Kahneman, D. (1992). Experiences extended across time: Evaluation of moments and episodes. *Journal of Behavioral Decision Making*, *5*, 169–185.
- von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton, NJ: Princeton University Press.
- von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research*. Cambridge, England: Cambridge University Press.

Received November 30, 1992

Revision received September 22, 1993

Accepted September 23, 1993 ■