Agricultural Economics Report

DECISION ANALYSIS IN AGRICULTURAL SETTINGS: AN INTRODUCTION

By

Lindon J. Robison
Beverly Fleisher

Department of
Agricultural Economics
MICHIGAN STATE
UNIVERSITY
East Lansing

DECISION ANALYSIS IN AGRICULTURAL SETTINGS:

AN INTRODUCTION

by

Lindon J. Robison
and
Beverly Fleisher*

ACKNOWLEDGMENTS

TABLE OF CONTENTS

ii

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

AN INTRODUCTION AND OVERVIEW

A decision maker decides between alternatives which we call action choices. If there are at least two action choices and the relative attractiveness of the choices is not clear, the decision maker has a decision problem. A decision problem exists, in most cases, because the action choices have multiple outcomes which must be compared. Converting the outcomes of action choices to common dimension, a process called premaximization, is the first step in resolving decision problems.

In many cases decision problems arise because outcomes from action choices are random variables. When action choice outcomes are not known with certainty we describe them by stating their likelihood of occurrence. Choosing between action choices with uncertain outcomes is also a decision problem, a problem so common that a science of decision making, decision theory, has emerged to address the problem.

The decision problem existing under uncertainty requires the same approach as does the solution to the decision problem when it exists under certainty. Each action choice must be reduced to a single dimensional index number compatible for comparison. By reducing each action choice to a single index number, reflecting its preference ranking, the decision problem can be solved. Rules designed to achieve this reduction we call decision rules. Describing decision rules is the goal of Chapter 3 in this volume.

Unfortunately, with any emerging discipline such as decision theory, the first task is to develop a vocabulary for communication between scientists. Although working on a common problem, decision scientists often find their academic homes in other disciplines. Attempting to bridge the communication barrier

1

between scientists and researchers with diverse backgrounds is the task addressed in Chapter 2.

The consensus of opinion, for the moment at least, appears to be that the expected utility hypothesis (EUH) is the most useful and reliable predictor of decision making behavior. Its usefulness as a theoretical tool as well as an applied tool has entrenched the EUH as the cornerstone of the science we call decision theory. Nevertheless the detractors of the EUH have been relentless. In some cases they have offered improvements. In some cases they offer evidence about the violation of the assumption underlying the hypothesis of the EUH. In some cases they present paradoxes not easily resolved by application of the EUH.

Chapter 4 is an evaluation of the EUH. It summarizes our experience with it as a predictive tool and discusses how the evidence supports or detracts from the hypothesis. This chapter draws heavily upon an article by the senior author of this volume which appeared in the American Journal of Agricultural Economics. We thank the editor of that journal for permission to use the material of that article in this report.

Scientific inquiry leads naturally to measures and comparisons. The statement that an individual has risk averse attitudes or is risk averse is naturally followed by the question: is individual A more risk averse than individual B? Failure to answer that question will impede the progress of decision scientists since the subjective nature of risk attitudes will require personalized prescriptions to resolve decision problems.

In Chapter 5 we summarize the tools available for ordering individuals according to their risk preferences. The limitation of the measures are also addressed.

The desire and the need to order individuals according to their risk pre-ference extends to the desire and need to order the probability density functions (pdfs) of action choice outcomes. If we say action choice 1 is risky, can we say: action choice 1 is more (less) risky than action choice 2?

Answering the above question is the subject of Chapter 6. The answer to the question, whether action choice 1 is more (less) risky than action choice 2, depends on the attitudes towards risk of decision makers choosing between action choices 1 and 2. Establishing this link between decision makers and riskiness of the action choices is the main topic the reader of this volume should expect in Chapter 6.

Chapter 7 explores recently developed extensions to the EUH. These include such innovations as the derivation of the EUH with fuzzy sets, prospect theory, and the generalized EUH without the independence axiom.

UNCERTAINTY, RISK, UTILITY FOR INCOME
AND ATTITUDES TOWARDS CHANCE TAKING

Introduction

This chapter should be viewed as the foundation for the remainder of the report. Its task is to provide the background for the discussion of decision makers' attitudes under uncertainty which will follow. Given our interest in understanding how decisions are made in the face of uncertainty, we must first review how decision makers describe their action choices which have uncertain outcomes.

To do this we introduce the subject of probabilities. This immediately creates a problem because there is wide disagreement about how probability measures originate. Fortunately there is unanimous agreement about their characteristics. These characteristics include: (1) the probability of an event occurring cannot be less than zero; and (2) the sum of probabilities of all possible mutually exclusive events must equal one. There are, of course, more elaborate rules for using probabilities but they needn't concern us in this discussion.

It is assumed in all decision models that action choice outcomes are described in probabilistic terms. In decision models, probabilities are considered primitives whose values are objectively given. But for the decision maker who must formulate probability estimates, such a cavalier dismissal of the probability measurement problem is not permitted. In forming probability estimates, the decision maker may find little comfort in the traditional objective probabilistic approaches.

The traditional schools of probabilities can be quickly summarized. The first approach centers around the notion of "equal likelihood". In this view an

experiment having n different outcomes which are assumed to have an equal chance of occurring is considered. To construct an experiment, m less than or equal to n of the possible outcomes are grouped together. Call the occurrence of a member of the group of m events E. The probability of event E is m/n. Defined in this way probability is simply a relative frequency measure for an event in an equally likely experiment. But elementary outcomes, particularly in agriculture, are not equally likely and, in most cases, the set of observations is not complete. For example, suppose we wish to measure the probability of monthly rainfall in a particular location. Most would argue rainfall follows cycles so that all possible rainfall observations in a given month may not be equally likely. Secondly, the question arises as to how many months of observations are needed to complete the data set. The answer is the data set is never completed because each year 12 new observations are added. The confounding of probability with its measure renders this approach to probability measurement almost meaningless for applied problems.

The second approach improves upon the first. It distinguishes between the concept of probability and its measure. Jacques Bernoulli defined probability as a "degree of confidence" for the occurrence of a particular event. The question becomes one of how confidence is established. One approach is to conduct experiments from which relative frequencies are measured. In this manner, the degree of confidence becomes the limiting value of the frequency of a favorable outcome (Venn). Schoemaker criticizes this approach for three reasons. First, probabilities taken as limiting values are never exact. Second, it is unclear what constitutes a replicable experiment from which data can be used to form "confidence". The outcome of a coin toss is often used as an example which produces data from which can be inferred probabilities of heads vs. tails in the limit. But if the experiment were exactly replicated, either heads or tails would always

appear. The uncertainty in the outcomes arises strictly because the experiment is not exactly replicated. And the lack of knowledge arises because the influence of the uncontrolled or inexactly replicated factors is not known. Some may argue, though, that the level of information required to know the outcome of a coin toss with certainty is unknowable (e.g., Heisenberg's uncertainty principle in physics) and that an acceptable replication is the control of known factors.

The third attempt to define probabilities is based on logic. This approach begins with a set of axioms and definitions (Jeffrey) which are consistent with intuitive notions of how the probability of an event may be determined. Then the consequences of the relationship between probabilities are mathematically deduced. The deduced relationships are used to test the truth of some hypothesis being questioned. In this instance, probability measures the logical, objective evidence assessed by a rational person.

In contrast to the objective probabilistic schools described above is the subjective or personal school of probability (Ramsey, de Finetti). In this view probabilities are degrees of belief subject to provisions of consistency which restrict the probability of events to be nonnegative and the sum of probability of mutually exclusive events in a universe to be one. This view allows, for example, the assignment of probability to a nonrepetitive event for which no information to make objective inferences is available.

Subjective probabilities do not necessarily exclude experimental or other data, but do not require it either. It is this relationship between experimental observations, data and subjectively held beliefs which cause problems in our interpretation of decision theory model results, particularly using the most important decision theory paradigm, the expected utility hypothesis. To illustrate the problem, consider this example. Suppose a probabilistic assess-

ment is being made about the amount of rainfall to expect in a particular location. The measure of interest is "inches of rainfall per year," denoted "r". Suppose variable measures have been kept for 50 years which can be used to form the objective probability density function $f(r)$ which measures the percentage of time during the 50 year period that annual rainfall equalled r for $0 < r < b$, where b is the highest rainfall observed. The subjective expectations of the decision maker need not equal $f(r)$ according to the subjective school of probabilistic measures. The event of interest, after all, is the likelihood of rainfall in a year not observed. So the subjective distribution, influenced by $f(r)$ can be written as:

$$g(r) = h(f(r)) \text{ for } 0 < r < b*$$

where h is some transformation of f and b* may differ from b.

Most decision theorists agree that $g(r)$ is the relevant probability measure to use when examining the actions of a decision maker. But suppose that no data on r are available, yet the uninformed subjective view is that the distribution is as before $g(r)$. Without the supporting evidence of $f(r)$, will the decision maker act exactly as before? The answer is "probably not". Yet most of the decision models ignore this fact. Knight made the distinction between the objective probability distribution $f(r)$ and the unsupported subjective probability distribution $g(r)$ which the decision maker may be unwilling to specify. The former he referred to as a risky choice while the latter was considered uncertain.

Modern decision theorists have dropped Knight's distinction because of the difficulty already discussed in the measurement of objective probabilities. Likewise, in most analyses, the confidence of the decision maker in the probability measure $g(r)$ is not discussed.

But Knight had a useful idea. Subjective probability must be uniquely defined because of rationality restrictions. In theory they are treated as perfect probabilistic measures of the choice environment. In practice they are differentiated by the confidence of the decision maker in his probabilistic information. So, the subjectively supported distribution $g(r)$ will likely be viewed differently than the unsupported distribution $g(r)$ despite the fact that they may be identical.

Unfortunately, decision theory has not learned how to incorporate this confidence factor into a decision model, despite the recognition of its importance. One approach for resolving this shortcoming was suggested by Meyer and Pope. They suggested that one could rely on sampling theory to reduce the dispersion of the probability density function as the number of observations of the elementary outcome increases. While this approach has some appeal, it also has a serious limitation. The limitation is that the probability function must be specified ex ante; usually it is assumed to be normally distributed. This unfortunately requires too much. We simply are not prepared to define the shape of the density function beforehand, especially in light of the fact that the empirical distribution is the unbiased probability density function measure.

So decision theorists remain in a dilemma. Without specifying the probability distribution ex ante, the number of observations supporting the distribution cannot be used to reflect the confidence in the probability measure. On the other hand, to assume a distribution ex ante violates the generally accepted view that probabilities of interest are subjective.

One possible solution is offered by Kahneman and Tversky who argue that subjective probability formulations may not follow consistency requirements. We might add the hypothesis that their adherence to such rules may be directly related to the objective data available. There does not appear, however, to be

reliable empirical data to support such a claim, apart from the fact that it resolves some heretofore unresolved decision paradoxes.

Now, having identified the probability measures in decision theory, we will ignore them in the remainder of this review. In this we follow the established approach. Probabilities will be treated as primitives whose values are determined by the decision maker who, we assume, holds them with complete confidence.

Risk Versus Uncertainty

Knight, in his seminal work, Risk, Uncertainty and Profit distinguished between risk and uncertainty on the basis of the amount of information available about the likelihood of outcomes of action choices. More specifically, risk required empirical information to generate probabilities. Uncertainty lacked this empirical base. However, the view that all probabilities are subjectively formed makes Knight's distinction irrelevant. Thus we find ourselves needing new definitions for risk and uncertainty.

Stiglitz appeared to be pessimistic about such definitional efforts. His contribution on the subject was that: "Risk is like love, we all know what it means but we can't define it." Despite such a pessimistic forecast we offer definitions for both "risk" and "uncertainty." These definitional efforts are described more completely in Robison and Fleisher. We draw heavily from our earlier work.

We begin our definitional efforts by introducing a primitive into our discussion. Our primitive, an undefined word, we call an "outcome." Synonymns for "outcome" include: events, happenings, response to an action choice, or results. Outcomes may be active or passive; they may occur as a result of a decision maker's actions, or independent of the decision maker. They may result in an improvement or a reduction in the decision maker's well being or leave him unaffected. They may be foreseen or unforeseen, or result in changes which are permanent or temporary. They are simply outcomes.

Uncertain and certain are <u>adjectives</u> used to describe outcomes. We cannot describe risk or uncertainty without first associating the concepts with outcomes. Uncertain outcomes are those with more than one possible outcome. To say an outcome is uncertain is to say that there exists more than one possible outcome with a positive probability assigned to its occurrence. Action choices with only one possible outcome are defined as certain; the single outcome has a probability of one of occurring. Since outcomes are either certain or uncertain, statements comparing the uncertainty of outcomes are inconsistent with our definition of the terms.

A class of uncertain outcomes which alter the well-being of either a well defined class of decision makers or a single decision maker are called "risky outcomes." "Riskiness," because it depends on the decision makers' attitudes and likes and dislikes, cannot be made more precise without first defining whose well-being is being used to give meaning to the concept. Once we define the class of decision makers, we may be able to make comparative statements like action choice A's outcomes are more (less) risky than B's. The important point to remember though is that an outcome's riskiness depends on the preferences of an individual or a class of individual decision makers. It cannot and should not be used interchangeably with the word uncertainty.

Our definition of risky outcomes comes close but does not correspond exactly with the ordinary use of the term risk. Part of the difference is due to the fact that we use "risky" and "uncertain" as adjectives to describe an outcome, not as nouns. "Risk" used as a noun, according to the dictionary, is the possibility of loss or injury. In this sense risk is used synonomously with the possibility of adverse outcomes. This definition is simply too restrictive for use in decision science; it must include the possibility of favorable as well as unfavorable outcomes.
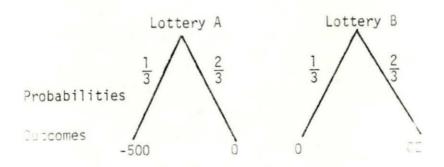
Utility of Income Versus Attitudes Towards Chance Taking

Risky action choices have uncertain outcomes whose occurrence alters the well-being of decision makers. Comparative statements about the riskiness of action choices then requires a statement or a description of the attitudes towards chance taking and preference for the outcome by the decision maker. During the past five years, researchers from several disciplines have been writing on the difference between utility of income, strength of preference, and risk aversion (Bell and Raiffa, Johnson, Kryzysztofowicz, Miyamoto, Sarin). We introduce this distinction with an example.

Suppose you face a risky event whose outcomes measured in dollars are either -$500 or zero. How much would you pay for an insurance policy which protected you against the possible loss of $500? The dollar figure given can be called the certainty equivalent, CE, for the risky event. From the expected utility hypothesis we could express your indifference between the insurance payment and the lottery as:

(2.1) $1/2\ U(-500) + 1/2\ U(0) = U(CE)$

Now answer the following question. Suppose you were faced with two risky events--of which you must choose one. One lottery has a 1/3 probability of obtaining -$500 and a 2/3 probability of obtaining zero. The second lottery has a probability of 1/3 of an outcome of zero and a probability of 2/3 of obtaining the certainty equivalent, CE, from the lottery A. Below we express these with tree diagrams.

Most individuals express a clear preference for lottery A or B despite the fact that the expected utilities of the two lotteries are equal.

(2.2) $1/3 \ U(-500) + 2/3 \ U(0) = 1/3 \ U(0) + 2/3 \ U(CE)$

A clear preference for lottery A or B is inconsistent with the earlier indifference between the lottery and the payment of the insurance premium. The inconsistency arises, we hypothesize, because of the preference for or the aversion to chance taking or gambling. This preference for or aversion to gambling or chance taking is separate from the utility of income which might be measured by ascertaining a level of income, y, at which the satisfaction gained from increasing from y_1 to y would equal to the satisfaction of increasing one's income from y to y_2 where $y_1 \leq y \leq y_2$. In other words, we can measure the simple utility for income by finding X such that:

(2.3) $U(y) - U(y_1) = U(y_2) - U(y)$

Arbitrarily assigning utility values to $U(y_1)$ and $U(y_2)$ allows us to solve for $U(y)$. Repeating the procedure allows the assignment of utility values to other income values. Note that this method for determining the utility of income involves no concommitant assignment of probabilities. Utility measures derived in this manner, therefore, involve no chance taking. The mathematical psychologists refer to utilty of income derived in this manner as a measure of strength of preference (Sarin).

The difference between utility functions derived using the technique described above rather than the certainty equivalent or the Ramsey method is reflected in their curvature. Miyamoto has hypothesized that utility functions derived without forcing the decision maker to take chances are more nearly linear than those derived using certainty equivalent or Ramsey methods.

The bending of the utility function measured by comparing risky choices does not necessarily imply an aversion to or a preference for chance taking. An

individual may have diminishing marginal utility and a preference for chance taking which, when combined, may result in either a concave or a convex utility function. As a result, we should not confuse the bending of the utility function with any particular attitude towards chance taking.

So what do our utility functions tell us? If they are derived by offering choices between a lottery and a certain income, the measure is a confounding of preferences for income and chance taking and is only a reliable measure in the absence of preference for or aversion to chance taking or in the comparison of two risky choices. This unusual state of neutrality toward chance taking is the basic assumption underlying the expected utility hypothesis.

Unfortunately we have become accustomed to using the terms risk averse, risk neutral and risk preferring to describe concave, linear, and convex utility functions. There is, we believe, little hope for correcting this misuse of the terms. However, we should recognize that preference for, aversion to, or neutrality towards risk or chance taking are attitudes distinct from the utility measure of income.

To be consistent with established use of the terms risk averse, risk neutral and risk preferring we propose the following. Decision makers whose utility functions are obtained by offering choices between lotteries are risk averse, risk neutral or risk preferrers over the income ranges which their utility functions are respectively concave, linear or convex. Requiring that the utility function be measured using comparison between lotteries assures us that attitudes towards chance taking as well as preference for income are included in the measure, and that the bending of the utility function is a composite measure of the two influences. This definition does not, however, allow us to determine the respective effects of preference for income and attitude towards chance taking.

To separate definitionally the influences of attitudes towards chance taking and preference for riskless income y, we follow Krzysztofowicz and refer to measures of the latter as value functions, $v(y)$, defined over income y. The mapping $w(v(y))$, which accounts for attitudes towards chance taking as well as preference for income, we refer to as the utility function $u(y)$ where

(2.4) $u(y) = w(v(y))$

where w is an individual's attitude towards chance taking.

Riskiness Versus Preference

Our lack of distinction between strength of preference for income and attitudes towards chance taking has lead to a related misuse of the words riskiness and preferences.

Earlier we defined risky events as uncertain ones whose outcomes may alter the well-being of decision makers. A problem arises when we make statements such as action choice A is riskier than action choice B and infer that A is less preferred than B. This implies not only an ordering of action choices according to preference, but also an ordering of decision makers according to their aversion to chance taking.

To make the distinction between uncertain and risky events more precise consider the following example. Decision makers 1 and 2 are considering action choices A and B whose outcomes are measured in dollars and have expected values E and variances σ^2 of (E_A, σ_A^2) and (E_B, σ_B^2) respectively for action choices A and B. Moreover, let $E_A > E_B$ and $\sigma_A^2 > \sigma_B^2$. Therefore according to our definition, events associated with action choices A and B are uncertain. Moreover, since both decision makers' well-being will be altered by the outcomes the action choice outcomes are also risky. Now since $\sigma_A^2 > \sigma_B^2$ can we say A is riskier than B? What if for individuals 1 and 2, who both have diminishing marginal utility for income and are adverse to risk taking, individual 1 prefers A and individual

2 prefers B? Then if riskier is to be used synonomously with "preference for" by "risk averse" decision makers variance is not an appropriate risk measure; it will only become such for decision makers whose utility function is a concave down quadratic function and whose distribution of events associated with action choices have an economic expectation.

Thus, measures of riskiness which imply "preference for" by "risk averse" decision makers are obtained only after specifying the class of decision makers for whom the measure applies. We also now recognize that orderings of action choices or decision makers must be based on a composite functional measure of attitudes towards chance taking and preference for income. We will return to this subject in Chapter 5.

Summary

Our definitional base has now been established. Probability measures are subjective. Thus the distinction between risk and uncertainty which depends on the distinction between empirically derived and distributions without an empirical base is largely rejected. In its place we propose that when there is more than one possible outcome, we refer to them as uncertain. When these outcomes may alter the well-being of decision makers, we describe them as risky. Risky outcomes then are a subset of uncertain outcomes.

Action choices facing decision makers are of interest when they have risky outcomes. However, whether one action choice is riskier than another by the nature of its outcomes depends on the utility of income and attitudes towards chance taking of the decision makers.

Unfortunately some confusion has arisen in decision theory because scientists have not distinguished between the two. We call the measures of preference for riskless income value functions. The function which is a composite of the value function and attitudes towards chance taking we refer to as a utility function.

Finally, to be consistent with decision scientists, we describe individuals with concave utility functions (who may either prefer or be averse to chance taking) as risk averse. Similarly, decision makers with convex utility functions (who may either prefer or be averse to chance taking) as risk preferring.

While the descriptors applied to the utility functions are not clear, the confusion introduced by altering them would, we hear, be worse. So we proceed.

CHAPTER 3

A REVIEW OF DECISION MODELS AND DECISION MAKING RULES UNDER RISK

Introduction

We previously defined an action choice as having uncertain outcomes if more than one outcome was possible. Decision theorists are interested in describing and comparing the outcomes of action choices with risky outcomes. As is customary in decision theory, the outcomes of interest are often expressed in terms of net return to the decision maker. This measure, regardless of the unit of denomination, is a net figure representing what is left for consumption after meeting obligations. The decision maker is assumed to have more than one action choice available which can be denoted as as a_1, a_2, ..., a_n. The outcomes which may result from an action choice depend on unknown or random states of nature denoted s_1, ..., s_m which the decision maker assigns probability measure $g(s_i)(i=1, ..., m)$. Consistency requires that $g(s_i)$ be nonnegative, and $g(s_i) + g(s_2) + ... + g(s_m)$ equal one. The interaction of the action choice of the decision maker and the possible states of nature is described as $0_{i,j}$ ($i=1$, ..., m; $j=1$, ..., n) where $0_{i,j}$ is the outcome resulting from the occurrence of the i-th state of nature given the decision maker's choice of the j-th action choice.

The elementary outcomes $0_{i,j}$ may be in nonhomogeneous units. For example, $0_{i,j}$ may be yields of soybeans per hectare, while $0_{i,k}$ may be yields of corn per hectare. With such nonhomogeneous measures, the outcomes must be converted first to a homogeneous measure. One approach is to measure the $0_{i,j}$'s in terms of their exchange for money equivalent, a measure $y_{i,j}$ where $y_{i,j}$ is the cash value of the outcome resulting from the j-th action choice occurring in the i-th state of nature. This conversion to a homogeneous measure, is related to what Johnson describes as "premaximization."

TABLE 3.1

A TABULAR DESCRIPTION OF A DECISION ENVIRONMENT INCLUDING ACTION CHOICES a_j (j=1, ..., n) STATES OF NATURE s_i (i=1, ..., m) and PREMAXIMIZED OUTCOMES $y_{i,j}$

States of Nature	Probability of Nature States	Action Choices		
		a_1	a_j	a_n
		Action Choice Outcomes		
s_1	$g(s_1)$	$y_{1,1}$		. $y_{1,n}$
.	.	.		.
.	.	.		.
.	.	.		.
s_i	$g(s_i)$	.	$y_{i,j}$	.
.	.	.		.
.	.	.		.
.	.	.		.
s_m	$g(s_m)$	$y_{m,1}$		. $y_{m,n}$
Probability density functions defined over ascending values of $y_{i,j}$		$g_1(y)$		. $g_n(y)$
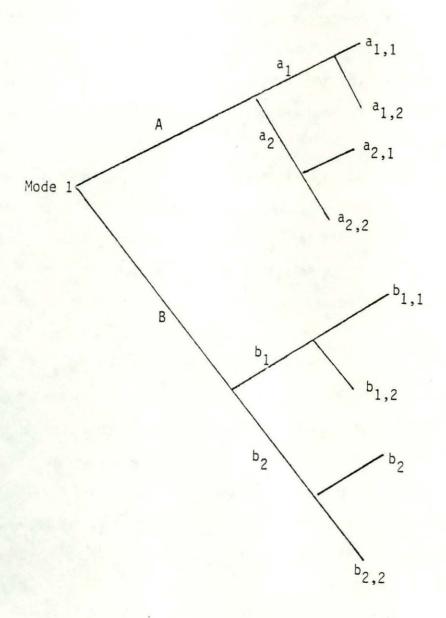
FIGURE 3.1

A GRAPHICAL DESCRIPTION OF A SEQUENTIAL DECISION
PROCESS WHICH BEGINS WITH A CHOICE BETWEEN A AND B
AND DEPENDING ON THE OUTCOME RESULTS IN A SECOND
CHOICE BETWEEN $a_{1,1}$, $a_{1,2}$, $a_{2,1}$, $a_{2,2}$, $b_{1,1}$, $b_{1,2}$, $b_{2,1}$ and $b_{2,2}$

Table 3.1 describes this decision environment. The first column lists the possible states of nature while the second denotes the subjective probability of each state's occurrence. The next n columns designate the action choices available to the decision maker. The premaximized outcomes $y_{i,j}$ in the body of the table define the interaction between an action choice and the occurrence of a state of nature.

The simplicity of the decision environment described in Table 3.1 should not lead one to conclude that it is irrelevant to complicated decision processes. The outcomes may result from sequential decisions, multiperiod outcomes, or have multiple attributes. Whatever the nature of decision environment, the choices eventually reduce to action choices described in terms of their probability density functions. Consider, for example, a sequential decision problem described by the decision tree in Figure 3.1. The decision maker begins with a choice beteen branches A or B. The possible outcomes of choosing A are a_1 or a_2 while the possible outcomes of choosing B are b_1 or b_2. If the deciion maker chooses A and a_1 occurs, then the decision maker may choose either $a_{1,1}$ or $a_{1,2}$ each with its unique probabilisitic set of outcomes described by $f(a_{1,1}/a_1)$ (read the probability of outcomes from action choice $a_{1,1}$ given that a_1 has occurred). Each branch constitutes a separate action choice, which depends on the intermediate outcomes of the first choice between A and B.

Outcomes occurring over time can be treated using similar approaches. One method is simply to find the present worth of the outcome received in the future. This is accomplished in the premaximization process.

Despite complications, the decision problem eventually collapses to comparison of actions a_1, ..., a_n described by their respective subjective probability density functions $g_1(y)$, ..., $g_n(y)$. The comparisons must be based on the shape of these probability density functions. Measures of the distribution include the

range of possible outcomes, the most likely outcome, or the mode, the center of weight or mass of the function, referred to as the mean or, more generally, the distribution's expected value. Another popular measure is the average deviation of outcomes from the mean, the standard deviation, or its square, the variance. All of these measures are helpful in describing and comparing probability distributions of action choices. The two most commonly used measures in decision theory though, are the expected value and variance of a distribution.

The expected value of the j-th action choice a_j is:

$$(3.1) \quad E(a_j) = y_{1j} \, g_j(y_{1j}) + y_{2,j} \, g_j(y_{2,j}) + \ldots + y_{m,j} \, g_j(y_{m,j})$$

When $g_j(y_{i,j})$ equals $1/m$, the expected value $E(a_j)$ is referred to as the mean.

To find how much, on average, an outcome varies from the expected value of a distribution, we compute the variance of the action choice denoted $\sigma^2(a_j)$ equal to:

$$(3.2) \quad \sigma^2(a_j) = f_j(y_{i,j})(y_{i,j} - E(a_j))^2 + \ldots + f_j(y_{m,j})(y_{m,j} - E(a_j))^2.$$

The squared deviations from the expected value provide an average squared distance of the outcome from $E(a_j)$. Taking the square root of $\sigma^2(a_j)$ provides the average difference of an outcome from its mean.

A long tradition has designated that if two action choices a_k and a_j have the same expected value, but the k-th action choice has a larger variance, it is described as "riskier" or, in the terminology of this paper, it is more random (Markowitz, Tsiang, Baumol).

It seems acceptable to refer to a distribution as more random than another if the probability weight is more dispersed. Certainly a distribution whose outcomes deviate a greater distance from its expected value than another with the same distribution qualify as being more dispersed.

Rothschild and Stiglitz point out that an increase in dispersion can result from shifting probability from the center of a distribution to its tail or by

adding a random variable with a zero mean to the original distribution. Both of these actions will increase the variance of the probability distributions and seem consistent with our notions of increasing randomness or risk.

Now we again enter murky definitional waters. We have discussed how a probability distribution becomes more random. We equated it with an increased variance or a spread in the probability function, leaving the expected value of the distribution unchanged. It is not uncommon to have distributions compared based on their riskiness. In fact, the mean preserving spread notion has been used to order action choices according to their riskiness. But riskiness of probability distributions connotes preference and we cannot indicate a preference ordering unless we have a preference ordering rule. Moreover, any ordering of probability distributions will likely vary between individuals precluding a complete and general order of probability density functions.

So the task at hand appears to be a review of decision rules that describe individuals' choice behavior under uncertainty. This theory as a minimum should allow us to order individuals according to their attitudes toward risk. Moreover, once individuals have been described according to their attitudes toward risk, we may use this information to order action choices according to their "riskiness" or preferability. For this chapter, however, we are content to summarize some of the decision rules (and the theories from which they are sometimes obtained) which reduce action choices to a single index of preferences.

The Decision Problem with Uncertain Outcomes

We have discussed methods for describing action choices. In what follows we discuss alternative approaches for indexing action choices so they can be ordered according to preference. This chapter begins with the simplest of rules and concludes with the expected utility hypothesis which Schoemaker describes as "the most generally accepted decision paradigm." We could also add that it is

TABLE 3.2

A COMPARISON BETWEEN THE i-th AND j-th ACTION CHOICE

States of Nature	Probability of Nature States	Action Choice a_i	a_j
		Premaximized Outcomes	
s_1	$g(s_1)$	$y_{1,i}$	$y_{1,j}$
.		.	.
.		.	.
.		.	.
.		.	.
.		.	.
s_m	$g(s_m)$	$y_{m,i}$	$y_{m,j}$

the basis for almost all of the disciplinary work being done in the economics of uncertainty.

A useful starting point is to review the decision problem introduced by uncertainty. A decision problem exists when the possible consequences of a decision are important and the best choice is not obvious (Anderson, Dillon, and Hardaker). So we might begin by decribing a decision in which the choice is obvious.

Using the language of Chapter 2, suppose a decision maker is faced with action choices a_1, . . . , a_n whose outcomes $y_{1,1}$, . . . , $y_{1,n}$ occur with probability of one in state one. In addition, assume that for units of income y, more is preferred to less. As a consequence, if $y_{1,j}$ is greater than $y_{1,i}$, it is preferred. The obvious choice among action choices then depends on the magnitudes of $y_{1,1}$, . . . , $y_{1,n}$ with the largest being preferred. The value of y then serves as an index which can be used to infer a preference ordering. We might, if we choose, transform the values of $y_{1,1}$, . . . , $y_{1,n}$ by a function U to create a new index $U(y_{1,j})$ ($j=1$, . . . , n) and the preference ordering would be unaffected as long as the function U were a monotonically (always) increasing function of y. As a result it makes little difference if we maximize the function $U(y)$ or the values of y to find the preferred action choice. They both produce the same result. The traditional approach of economists has been to ignore the function $U(y)$ and maximize over y.

Now consider the complication introduced by uncertainty. With the introduction of uncertainty, the comparison between action choices has been complicated because of the multiplicity of outcomes which may occur with probability greater than zero. Consider, for example, the pairwise comparison between the i-th and j-th choices described in Table 3.2. Since there are m possible outcomes under each action choice there are m^2 possible comparisons between

outcomes of the two action choices above. And if all pairwise comparisons of n action choices each with m possible outcomes were made, the number of comparisons N equals:

$$N = (n)(n-1) \ . \ . \ . \ (n-3)m^2$$

a number very large even for reasonably small values of m and n. Moreover, the comparisons are of little value unless an indexing rule is available to rate differences in outcomes.

There is one case, of course, where the ordering is obvious even under uncertainty. Suppose $y_{k,i} \geq y_{k,j}$ for k=1, . . . , m. Then no matter which state of nature occurs, the i-th action choice has the most favorable outcome and is preferred. This condition meets Hadar and Russell's first degree stochastic dominance requirement. But this is a strong requirement, a requirement not likely to be met in most comparisons. For those choices where the inequality between $y_{k,i}$ and $y_{k,j}$ are reversed over at least one of the states of nature, preference will be unclear.

The comparison problem just described has led to rules which provide single value indexes. The number of such indexing rules is large; only a small sample of the rules will be discussed.

Maximax and Minimax Rules

The maximax indexing rule uses the maximum outcome which occurs under each action choice as an index. It searches the outcomes under each action choice for the maximum or the most favorable event. Suppose in Table 3.1 this is the outcome $y_{1,i}$ for the i-th action choice and $y_{1,j}$ for the j-th action choice. The values of $y_{1,i}$ and $y_{1,j}$ become the index values for the action choice and indicate preference. If $y_{1,j} \geq y_{1,i}$ for example, the j-th action choice is preferred.

The closely aligned alternative to the maximax rule is the minimax rule. Instead of focusing on the most favorable outcome, it focuses on the least favorable outcome. The index value becomes the worst that can occur. The largest outcome of the worst possible, is, of course, preferred. The action choice corresponding to the best of the worst outcomes is preferred.

The maximax and the minimax indexing rules describe the extremes of response to uncertainty. The maximax rule which is based solely on the most favorable outcome while ignoring all other possibilities reflects extreme optimism. In contrast, the minimax rule which focuses on the least favorable outcome is pure pessimism. To ignore the other possibilities and the probabilities with which they may occur is certainly an incomplete evaluative criterion.

It follows that alternative rules could improve upon the maximax and minimax rules by accounting for the probabilities and values of outcomes of alternative outcomes. In the process these rules could possibly capture types of behavior other than extreme optimism or pessimism. A step in this direction might be the mixed strategy model. This model attempts to provide an intermediate response to uncertain action choices. It does this by selecting an index, α, for each action choice.

The method identifies both the maximum or most favorable outcomes, $y_{max,i}$ and $y_{max,j}$, and the least favorable outcomes, $y_{min,i}$ and $y_{min,j}$, from the i-th and j-th action choices respectively. Then using α, a linear combination is formed equal to:

$$\alpha \, y_{max,i} + (1-\alpha) \, _{min,i} = y_i^*$$
$$\alpha \, y_{max,j} + (1-\alpha) \, _{min,j} = y_j^*$$

where y_i^* and y_j^* become the preference indexes for the action choices. The rule just described becomes operational, however, only when the decision maker can supply the coefficient.

The criticisms of this model are similar to those made of the minimax and maximax models. Why ignore all the values between y_{min} and y_{max}? And why don't probabilities matter? One response to these criticisms which applies to all of the models just described is that no data are available from which subjective probability density functions can be formed. As a result, the decision maker has no basis to infer anything about the distribution beyond its upper and lower values. In such a state of ignorance about the "true" shape of the probability density function, the decision maker is only justified in using upper and lower values in the decision rule used.

But if no data beyond high and low values is available, then each data point in between should be equally weighted which results in a uniform probability distribution shown in Figure 3.2. As a result, the models just described have little practical relevance, except to exemplify the extremes of optimism and pessimism.

Safety-First Models

An alternative to the maximax, the minimax and the mixed strategy model is the use of some version of the safety-first model. In its simplest form, the safety-first model focuses on a safety or disaster level of outcome y_d. This outcome may be an income level below which a firm fails to meet its cash obligations or becomes bankrupt. In a developing country setting, it may be the minimum level of returns required to satisfy survival requirements. Whatever the interpretation of y_d, this model assumes that its objective is to select action choices in such a way that chances of experiencing y_d or worse are minimized.

We illustrate this rule with action choices a_i and a_j by drawing their respective cumulative density functions in Figure 3.3. The cumulative density functions are obtained from their respective probability density function by summing. To illustrate, if $g_i(y_{k,i})$ is the probability density function of the
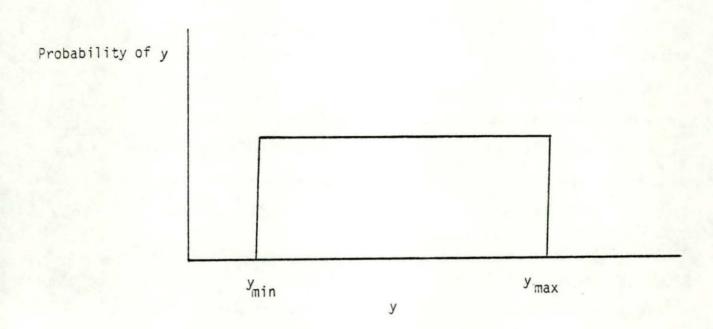
FIGURE 3.2

A UNIFORM PROBABILITY DENSITY FUNCTION IN WHICH EACH OUTCOME BETWEEN
THE MAXIMUM y_{max} AND THE MINIMUM y_{min} ARE EQUALLY LIKELY
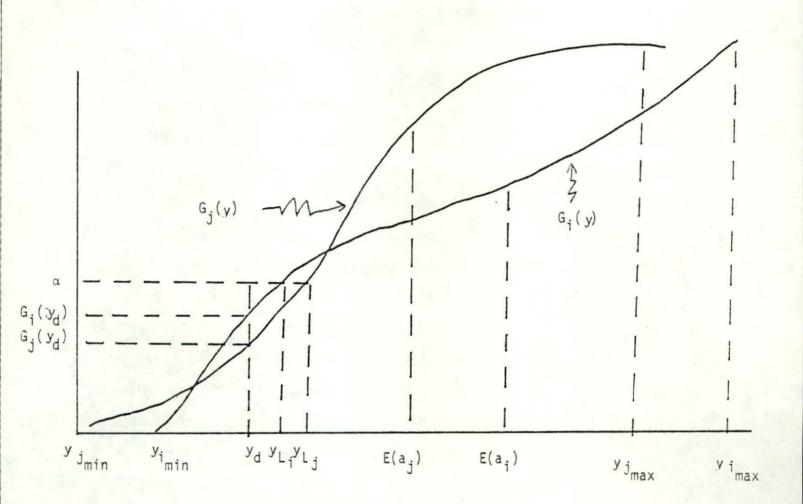
FIGURE 3.3

CUMULATIVE DENSITY FUNCTION $G_i(y)$ and $G_j(y)$ DESCRIBING
PROBABILISTIC OUTCOMES OF RECEIVING y_d OR SOMETHING LESS

i-th action choices, its cumulative density function values at $y_{k,i}$ is $G_i(y_{k,i})$ = $\sum\limits_{j=-\infty}^{k} g_i(y_{j,i})$ which obtains the sum of probability of outcomes equal to $y_{k,i}$ and below. The function $G_i(y_{k,i})$ can be read as the probability of $y_{k,i}$ or something less occurring. The maximum value $G_i(y_{k,i})$ can take on, of course, is limited to the sum of all probabilities of $y_{k,i}$ occurring, which is one.

Let $G_i(y_{k,i})$ and $G_j(y_{k,j})$ represent the cumulative density functions for the i-th and j-th action choices. These are represented in Figure 3.2. At the disaster outcome y_d, $G_i(y_d)$ exceeds $G_j(y_d)$ suggesting that the probability of y_d or something worse occurring is greater with the i-th action choice. Then, according to this version of the safety-first model, the j-th action choice is preferred, despite the fact that it has less favorable maximum possible outcomes $(y_{max,j} < y_{max,i})$ and a worst minimum outcome $(y_{min,j} < y_{min,j})$.

The safety-first model improves upon the earlier models by focusing on an outcome y_d which may be different than the worst possible or most favorable outcome. And if y_d marks the worst possible outcome for the decision maker, then outcomes below y_d may be safely ignored. But, why ignore the likelihood of outcomes more favorable than y_d? Shouldn't the distribution of probabilities of outcomes above y_d matter? The answer is yes. To ignore what happens above y_d is to assume that utility of $y > y_d$ is zero.

To account for outcomes above y_d in the context of a safety-first model, Pyle and Turnovsky suggested three alternatives. The first rule attributed to Telser assumes that a decision maker maximizes expected returns $E(a_k)(k=1, \ldots, n)$ subject to the constraint that the probability of a return less than or equal to a specified or disaster outcome y_d does not exceed a stipulated probability.

To illustrate this rule we return to Figure 3.3. Suppose the stipulated probability for experiencing y_d or something less is α. All action choices considered by the decision maker must have cumulative density function values at

y_d not greater than α. In Figure 3.3, both $G_i(y_d)$ and $G_j(y_d)$ are drawn so they meet the first constraint. Having satisfied this rule, the choice between the i-th and j-th action choices are made based on their respective expected values. In this case the i-th action choice is preferred.

A second safety-first rule proposed by Kataoka is based on a particular probability value of $G(y_L)$. It selects the action choice which maximizes y_L for given probability values of $G(y_L)$. In effect, this rule maximizes the minimum return which can be earned for a fixed value of $G(y_L)$.

To illustrate, let the fixed value of $G(y_L)$ or the lower confidence limit be α. Then the action choice with the largest value of y_L at a given value of $G(y_L)$ is preferred. In Figure 3.3, $G_j(y)$ is preferred to $G_i(y)$ since the value of $y_{L,j}$ exceeds $y_{L,i}$.

A third safety-first rule was developed by Roy. It identifies the optimal plan as one having the smallest probability of yielding a return below some specified level. This corresponds closely with the objective of avoiding a disaster level of income as the sole objective of the decision maker.

An Alternative View

While safety-first rules focus on probabilities or specified outcomes, an alternative approach considers all outcomes, given knowledge of the range of outcomes, y, and their likelihood function, $G(y)$. If each y influences the well-being of the decision maker, why should any be ignored in the decision rule? If indeed there is a disaster outcome whose occurrence has more impact on the decision maker's well-being than another, this could be accounted for by weighting it differently. As is most often the case, the simplest of these weighting rules was the first adopted.

The rule in general acceptance in the eighteenth century when mathematician Daniel Bernoulli was studying decision making was to weight outcomes according to

their monetary value. The index became the expected monetary value of each action choice and was used to rank the action choices of the decision maker.

This rule did what safety-first rules did not. It allowed each possible outcome to influence the preference index. This feature and its simplicity has made it a popular decision theory tool even today. One popular application has been in linear programming models. By replacing uncertain parameters in linear programs with their expected values, the outcome is the solution which maximizes the expected value.[1]

Bernoulli, however, found an inconsistency between the expected value rule and the way decision makers actually behave. The inconsistency arose in a gamble referred to as the St. Petersburg paradox. The gamble paid depending on the number of flips of a coin required to obtain heads. If, for example, heads occurred on the first flip, the gamble paid $2. If heads occurred on the second flip, the gamble paid $(\$2)^2$ or $4. And if heads occurred on the third flip, it paid $(\$2)^3$ or $8. The probability of heads occurring on the first flip is 1/2, 1/4 on the second flip, and 1/8 on the third flip and so on.

The expected value of gamble E(G) then could be written as the sum:

(3.3) $E(G) = 1/2 (\$2) + 1/4(\$4) + 1/8(\$8) + \ldots .$

The value of each individual element in the above gamble is one. The number of elements, however, is infinitely long so that sum, or the expected value, of the gamble is not finite.

If decision makers played this gamble according to the expected value rule they should pay any "large" amount to play since the expected value of the gamble is infinite. But as Bernoulli observed, the amount decision makers were actually

---

[1]For a more modern criticism of this approach, see Stovall.

willing to play was a finite (small) amount. This observation led Bernoulli to conclude that decision makers maximized some function other than the expected value.[2]

He proposed they maximize the log function of the premaximizing outcomes. This has become known as the Kelly criterion and is still popular today. The Kelly criterion, it turns out, is equivalent to maximizing the geometric mean of a gamble. And as others have pointed out, maximizing the geometric mean will maximize either the expected value of terminal wealth or minimize the number of plays required to achieve some wealth in a repeated gamble (Bierman).

The Kelly criterion does, however, have its shortcomings. Negatively valued outcomes are not defined by a log transformation. As result, an additional transformation would be required in order to maximize the expected log when negative outcomes are involved. A second question which arises is, "Do decision makers respond in identical ways, i.e., by maximizing the log or any other function?" While this is an empirical question, the evidence appears to indicate they do not (Love and Robison).

Still, the recognition that additional units of income may not be valued by a constant amount was a step forward. The concept of diminishing marginal utility of income corresponding to the log utility function was consistent with the output responses to inputs with which physical scientists were well acquainted. For this contribution, Bernoulli is credited with the notion of utility of income.

---

[2]Samuelson has questioned Bernoulli's conclusion. He asked: who has the infinite sum required to pay the gambler should the most favorable occur? Since the payoff is not available, no one should pay the large price to play.

But the assumption of utility of income being mapped by the log function was too strong an assumption to make. It remained for Ramsey and later von Neumann and Morgenstern to introduce the generalized expected utility theory which remains the most generally accepted decision paradigm.

The Expected Utility Hypothesis

The expected utility hypothesis (EUH) asserts that if a decision maker's behavior is consistent with a set of axioms which we will describe, they will weight outcomes according to a personalized (possibly unique) function $U(y)$. The expected value of $U(y)$ provides the single-valued index which orders action choices according to the preferences or attitudes of the decision maker.

A complete development of the EUH is found in the landmark work of von Neumann and Morgenstern and reviewed in Luce and Raiffa. Only the highlights of the theory are presented here.

The EUH assumes decision makers obey the following axioms and initial conditions. The initial conditions are that the decision makers can identify a set of action choices a_1, . . . , a_n and can associate with the action choices probability density functions $g_1(y)$, . . . , $g_n(y)$ respectively. The probability density functions are subjective and are assumed to obey the calculus of probability. The axioms of behavior fundamental to the EUH include the following:

1. Ordering of Action Choices

    For any two action choices, a_1 and a_2, the decision maker either prefers a_2 to a_1, or is indifferent between them.

2. Transitivity Among Choices

    If a_1 is preferred to a_2, and a_2 is preferred to a_3, then a_1 must be preferred to a_3.

3. Substitution Among Choices

If a_1 is preferred to a_2, and a_3 is some other choice, risky choice $pa_1 + (1-p)a_3$ is preferred to another risky choice $pa_2 + (1-p)a_3$, where p is the likelihood of occurrence.

4. Certainty Equivalent Among Choices

If a_1 is preferred to a_2, and a_2 is preferred to a_3, then some probability p exists so that the decision maker is indifferent between a_2 or receiving a_1 with probability p and a_3 with probability $(1-p)$. Thus a_2 is equivalent in satisfaction to the compound lottery $pa_1 + (1-p)a_3$. If a_2 is a single value whose probability of occurrence is one, it is an income received with certainty and is equivalent in satisfaction to the lottery. In this context, a_2 is referred to as a certainty equivalent income.

If a decision maker obeys these axioms, a utility function $U(y)$ can be formulated which reflects the preferences of the decision maker (Hey). A discussion of the procedures for measuring $U(y)$ follows while an overview of how well they predict decision maker's actual choices is reserved for Chapter 4. Inferences which can be made from the measurements of $U(y)$ is the subject of Chapter 5.

Measuring Decision Makers' Utility Functions

If each individual has a unique weighting of outcomes, then prescription of an action choice requires that the unique preference function $U(y)$ be measured. But before proceeding to that subject, we might ask why we need to measure the preference function at all. Why not present the decision maker with the action choice set and let him make a selection based on internal and unrevealed preferences? This, of course, is the procedure followed in the largest number of decision making processes.

But there are other decision situations where the action choice set is so large as to preclude its careful evalution by a decision maker. For example, consider an investment in a mutual fund. Currently there are 500 mutual funds exclusive of money market funds. Rather than make an effort to learn all about the funds himself, the decision maker may consult a broker. But the broker must have some information about the decision maker's investment preferences before a recommendation could be made.

In another case the action choice set may be the result of a computer simulation which creates a large number of possible action choices. In the case of Monte Carlo simulations, some screening device is required before a set of choices small enough for evalution by the decision maker is available.

In both cases, some characterization of the decision maker's preference is required before the analysis can proceed. For these reasons and others a description of the investment attitudes of the decision maker are required. We measure the function $U(y)$ for each decision maker by means of the expected utility hypothesis (EUH).

For a complete discussion of how the measurement procedure occurs, the reader is referred to an excellent technical guide by Halter and Dean, Chapter 4 in a more recent text by Anderson, Dillon and Hardaker, or a recent article by Halter and Mason.

The measurement of $U(y)$ begins by assuming decision makers can identify the most and least favorable outcomes y_H and y_L, respectively. To measure $U(y)$ over the range of y_H and y_L, the analyst constructs hypothetical gambles using various values of y between y_L and y_H and probability value p. All of the procedures seek to find an indifference between two lotteries or a lottery and a sure outcome by either adjusting the value of the outcome or the probability of occurrence. Finding this indifference supplies one new piece of information about the decision maker.

The procedure which adjusts probability to find the indifference point proceeds as follows. Assign y_L the value $U(y_L)$ and y_H the value $U(y_H)$. This is permitted as long as $U(y_H) > U(y_L)$; i.e., the ordering of y_H and y_L is not reversed. Then choose an intermediate value of y_m such that $y_L < y_m < y_H$. Since y_H is preferred to y_m and y_m is preferred to y_L, then according to the certainty equivalent axiom, there is a probability p such that the decision maker is indifferent between the lottery with expected value $py_L + (1-p)y_H$ and y_m received with certainty. The probability value p which makes the two choices equal in preferences is supplied by the decision maker.

The utility value of y_m, $U(y_m)$, which indicates indifference can be solved from the expression:

(3.4) $U(y_m) = pU(y_L) + (1-p)U(y_H)$.

Now three points on the function $U(y)$ are known $U(y_L)$, $U(y_H)$ and $U(y_m)$. But others can be obtained by forming a new lottery out of y_H and y_m or y_L and y_m. To illustrate, if y_H were $1,000, y_L were $0, and y_m were $400, we might begin by assigning $1,000 the utility weight of 100 and $0 the utility weight of zero. If the indifference probability supplied by the decision maker were p=.4, then the utility weight assigned to $400, U($400) is equal to 60:

(3.5) $U(\$400) = (.4)(0) + (.6)100$

$= 60.$

This approach, referred to as the von Neumann-Morgenstern model by Officer and Halter, has particular application for decision problems with discrete outcomes or outcomes not measured in the same units but which can be ordered. In such a case the probabilities, not the outcomes, are adjusted to indifference.

The modified von Neumann-Morgenstern is a similar model except that it adjusts the outcome to find indifference. In this model we search for an outcome y_{CE} which, if received with certainty, equals the lottery whose expected value is

$1/2 (y_L + y_H)$. The utility value assigned to y_{CE} can be found from the expression:

$$U(y_{CE}) = 1/2 (U(y_L) + U(y_H)).$$

To illustrate, if y_{CE} were \$350 and $U(y_L)$ and $U(y_H)$ were the values assigned previously, then

$$U(y_{CE}) = 1/2(0+100)$$
$$= 50.$$

The difficulty with both of the approaches is that they assume the decision maker is neither averse to nor prefers chance taking. Otherwise attitudes towards chance taking are confounded with utility of income using the von Neumann-Morgenstern or the modified von Neumann-Morgentern approach.

There are, of course, more complicated lotteries which could be constructed to find indifference. If one wishes to compare two uncertain lotteries, the Ramsey method is suggested. It starts with outcomes $y_L < y_1 < y_H$ and solves for the probability p which makes the two lotteries below equal in preference, that is,

$$1/2(U(y_L) + U(y_H)) = pU(y_L) + (1-p)U(y_1).$$

Knowing that the value of $1/2(U(y_L) + U(y_H))$ and p allows us to write:

$$U(y_1) = (1/2(U(y_L)) + U(y_H) - pU(y_L))/1-p.$$

Now three points have utility associated with them $U(y_L)$, $U(y_H)$ and $U(y_1)$. This allows us to construct still other lotteries and associate utility with still other points.

When enough utility values are available, a utility function can be fitted to these points using either graphical or statistical procedures. The statistical approach consists of postulating a functional form for utility and regressing it over the data points. Some error will be introduced in this process because the function selected will not exactly match the data. Nevertheless, the

expected value of the estimated function will serve as a single valued index for ordering action choices according to preference.

The utility measurement procedures all consist of searches for indifference. To find the indifference points, flexibility in the construction of action choices is required. As a result, the action choices are always hypothetical in nature, although the analyst may try to relate the hypothetical outcomes to real world events with which the decision maker is familiar. Since the utility function is usually derived for applications in other settings, it remains an exercise in mind experiments. Whether or not the utility information obtained from such experiments has applicability for real world decisions is an issue not fully resolved, although Binswanger claims his Indian study supports such a claim.

Properties of the Function U(y)

There has been some confusion in the past over whether the function $U(y)$ measured using one of the methods described is a cardinal or an ordinal function. The basic difference between an ordinal function and a cardinal function is the following. An ordinal function orders outcomes according to preference. Given outcomes $y_1 < y_2 < y_3 < y_4$, a monotonic ordinal function would have values such that $U(y_1) < U(y_2) < U(y_3) < U(y_4)$. A cardinal function $U^*(y)$ could do more. If $U^*(y_2) - U^*(y_1) < U^*(y_4) - U^*(y_3)$ we could infer that the additional satisfaction of increasing one's income from y_3 to y_4 was greater than going from an income of y_1 to y_2. This could not be inferred from the von Neumann-Morgenstern function $U(y)$ (Schoemaker)[3].

---

[3]An axiomatic system which allows for orderings on differences has been provided by Debreu.

Since the origins and scale are assigned arbitrarily when defining U(y), this ordinal function is not unique. Any positive linear transformation would do just as well. For example, suppose in the comparison of action choices a_i and a_j that their respective probability density functions were preference ordered as follows:

$$\sum_y U(y)g_j(y) \qquad \sum_y U(y)g_i(y).$$

The same preference ordering would be preserved by any function $U^*(y)$ equal to $a+bU(y)$ where a and b are arbitrary constants and $b>0$. This can easily be shown from the expression

$$\sum_y (a+bU(y))g_j(y) > \sum_y (a+bU(y))g_i(y)$$

which simplifies to

$$a\sum_y g_j(y) + b\sum_y U(y)g_j(y) \qquad a\sum_y g_i(y) + b\sum_y U(y)g_i(y).$$

Since $\sum_j g_j(y)$ and $\sum g_i(y)$ equal one, we can subtract the a's from both sides, cancel the positive parameter b and obtain the original inequality.

Concerns About the EUH

Despite the prominence of the EUH as a decision tool under uncertainty, it does have its detractors. They have raised questions such as: Are decision makers' true tendencies revealed in a game-like setting?, Are preferences over time constant?, And while all theories of behavior only approximate actual real world behavior, can a theory which includes income as its only independent agrument be accurate enough to be useful? Certainly the precision assumed in the EUH is unwarranted. Neither probability density functions nor utility functions can be measured without error. Therefore, we are left with an empirical question: How accurate is the EUH in applied settings? The next chapter reviews the evidence.

CHAPTER 4

EVALUATING THE PREDICTIVE ABILITY OF
EXPECTED UTILITY MAXIMIZING MODELS[1]

Introduction

The expected utility hypothesis (EUH) is primarily a prescriptive tool. It suggests that if decision maker behavior conforms to certain axioms, they will maximize their well-being by selecting action choices which maximize expected utility. We now ask whether decision maker behavior is consistent with the axioms. And if so, can the EUH be used as a predictive tool? We proceed to evaluate what we know about the predictive behavior of the EUH. But before doing that, we must carefully describe what conditions are required for a test of the EUH's predictive ability.

Constructing a Test of the Expected Utility Hypothesis

To claim that the EUH can be used to predict behavior is to claim that the EUH is a supported theoretical hypothesis; that is, we hypothesize that decision makers behave as if they were expected utility maximizers. According to Giere, support for a theoretical hypothesis requires an experiment or set of observations which involve the hypothesis, initial conditions, auxiliary assumptions, and a prediction. For the theoretical hypothesis to be supported, two conditions must be met: (1) if the auxiliary assumptions, initial conditions, and hypothesis are true, then a correct prediction will probably follow; and (2) if auxiliary assumptions and initial conditions are true and the hypothesis is not true, then the correct prediction will probably not be observed.

The word "probably" in the two conditions identifies the theoretical hypothesis as probabilistic rather than deterministic. The model will likely omit some features of the real world affecting decision making behavior. So we do not expect perfect prediction, only that the evidence does not permit a rejection of the model -- since truth of the model itself cannot be established.

Condition one requires an experiment involving the initial conditions, auxiliary assumptions, and hypothesis, which are used to make a prediction. The experiment's prediction is compared with actual decisions to determine if condition one is satisfied. To satisfy the second condition, the experiment's result must have a low probability of being predicted from an alternative hypothesis. If the same prediction results from many alternative hypotheses, the second condition would not be satisfied and the theoretical hypothesis is not fully supported.

Suppose we wish to examine the support for the theoretical hypothesis that decision makers order action choices according to the expected utility hypothesis. The hypothesis in such a test is the EUH. The initial conditions are the choice set with consequences described in probabilistic terms and the decision maker's utility function. The auxiliary assumptions are (a) that the decision maker's utility function and the probability density functions describing the consequences of the action choices are measured accurately and (b) that the axioms underlying the EUH are valid. Alternative theoretical hypotheses might be that decision makers order choices based on (a) expected profits, (b) safety levels of income, (c) lexicographic utility functions, or (d) expected losses. Finally, the prediction is the action choice selected by maximizing expected utility. Condition one requires that the prediction matches the decision maker's actual choice. Failure to predict the actual choice forces a rejection of either the hypothesis, initial conditions, or auxiliary assumptions. Obtaining the prediction from an alternative model causes the EUH model to fail condition two.

Describing the Choice Sets

To examine support for the EUH, two choice sets are required. One set is used to obtain the decision maker's utility function. From the second set, an expected utility-maximizing choice is predicted which is compared with the decision maker's actual choice.

At least two approaches can be used to construct choice sets. One approach is to describe the actual choice set facing an individual. This method is referred to as the actual economic behavior approach. For complex choice sets, however, the actual economic behavior approach is difficult and costly and therefore rarely used. As a result, researchers more often construct artificial choice sets. This approach is referred to as the experimental approach.[2] The experimental approach has been criticized because it forces the decision maker to respond to hypothetical questions. If the decision maker is an expected utility maximizer over all his resources, including his time, he may respond to hypothetical questions in such a way that minimize his cost (time) of participating rather than reflecting his preferences for the hypothetical outcomes.

To overcome this criticism, a third approach can be used. This approach, also an experimental one, satisfies the initial conditions by artificially constructing a choice set using significant outcomes and is referred to in this paper as the experimental approach with significant outcomes. The limitation of such an approach, of course, is that not all experiments can afford to reward respondents with significant outcomes. The exceptions would be in developing countries where significant outcomes may be small levels of income relative to the budget of the researchers.

---

[2]This definition of the experimental approach differs from Binswanger's. His definition of the experimental approach is described in this paper as the experimental approach with significant outcomes.

Identifying the Decision Maker's Utility Function

The remaining initial condition to be satisfied before proceeding with a test of the EUH is to identify the decision maker's utility function. Identifying the decision maker's utility function, however, requires that he be confronted with alternative action choices so that indifference can be established between alternative uncertain action choices or between a sure outcome and an uncertain action choice. Efforts to derive a decision maker's utility function have almost always relied on the experimental approach for defining the choice set, asking the decision maker to choose between hypothetical choices.

Several studies have assumed that decision makers' utility functions belong to a certain family of functions, usually ones that are described by a single parameter. Then, defining choice sets using either the experimental or actual economic behavior approach, they use the actual choice of the decision maker to solve for the parameter which identifies risk preferences. Examples of such an approach can be found in Brink and McCarl, Binswanger, and Dillon and Scandizzo, who all used an equilibrium slope on an EV set or an equivalent mean-variance trade-off measure: Binswanger, and Grisley and Kellogg, who used partial risk-aversion measures obtained for a specific gamble; and Dillon and Scandizzo, who used single parameters from an assumed quadratic or power utility function.

While this approach may provide useful information about the distribution of risk coefficients measured, inferring from a measured risk coefficient to a utility function is an unjustified approach. Moreover, the reliability of these coefficients in predicting expected utility-maximizing choices still needs to be established. And, until there is evidence to support the assumption that risk preferences can be inferred in this manner, the results of this type of study cannot be used to evaluate the support for the EUH.

The difficulty of meeting initial conditions aside, there have been satisfactory studies made of the EUH. They are organized below according to how the choice set was obtained. A separate class of studies reviewed examines the validity of the axioms from which the EUH is deduced. In the first study reviewed, the application of Giere's conditions is carefully considered. Because of space constraints there is less mention made of Giere's two conditions in the remaining studies.

Actual Economic Behavior Tests

Lin, Dean, and Moore constructed a comprehensive test of the EUH using the actual economic behavior approach to evaluate risk preferences on large-scale California farms. To describe the choices facing decision makers the authors constructed an expected value-variance (EV) efficient set for six farmers in the San Joaquin Valley. Utility functions were obtained using the experimental approach and predictions about farm organizations were made for each decision maker. The predictions resulted from maximizing expected utility for each decision maker over their respective choice sets. The test consisted of comparing the predictions with observed economic behavior. Condition two was satisfied by making predictions using expected profit and lexicographic models.

In only three of the six cases did the EUH model predict better than a lexicographic model or an expected profit-maximizing model. In none of the cases did the EUH model predict the actual farm plan followed by the decision maker. In fact, it would have been impossible for the EUH predictions to have made correct choices because the actual choices were not members of the choice set from which the predictions were selected. Thus, an important initial condition required for the test -- correct identification of the choice set -- was violated.

The authors then reconstructed the experiment and modified the initial conditions, restricting the decision makers' action choices to members of the EV set. In this approach, the authors adopt the experimental approach to define the choice set from which predictions are made. In the new test, the predictions matched the actual choices for three of the decision makers and came closer in the other three cases than the predictions from either the lexicographic model or the expected profit model. These results support the EUH.[3]

The Lin, Dean, and Moore study is appealing because both conditions one and two were used in the test. To meet condition one an experiment produced a test which followed from the EUH, initial conditions and auxiliary assumptions. Then condition two was examined to see if correct predictions could be obtained from alternative hypotheses. But as an application of the actual economic behavior approach, it failed to satisfy the auxiliary assumption that the choice set was accurately described because actual choices were not predicted nor included in the choice set. Brink and McCarl experienced similar difficulty in their attempts to model actual economic choices using the equivalent of a mean-variance set. This evidence suggests that, except in very simple decision environments, constructing a test of the EUH using actual economic behavior will be difficult.

A quite different lesson was learned from a test of the EUH made by Haneman and Farnsworth. They used the EUH model to test whether or not the EUH model could explain why one group of farmers adopted integrated pest management strategies (IPM) while another group continued with conventional (chemical) control programs. Using the experimental approach, they derive utility functions for both groups. They found no significant difference in the risk attitudes of the two groups; however, they did find significant difference in subjective

---

[3]Even stronger support for the EUH is implied by Lin and Chang. They increase the number of correct choices using the EUH model by estimating preferences using a more flexible econometric model to describe utility functions.

expectations on yields and profits between the IPM and chemical control groups despite the fact that historically there was no significant difference.

Their study results showed that for thirty-five of the forty-four decision makers either expected utility maximization or expected profit maximization predicted the pest control method actually being followed. And in five of the nine cases in which switching control strategies was recommended, the subjective probability density functions gave unclear preference signals.

While Haneman and Farnsworth produced a prediction consistent with the EUH and met condition one, their test failed to meet condition two because the same prediction was made by either the EUH model or the expected profit model. This evidence is weak support or lack of support for the EUH. Haneman and Farnsworth, however, infer something else. They infer that "subjective perceptions of outcomes rather than the type of choice criteria or the nature of risk preferences explain (the prediction)" (p. 19). The confirmation of this hypothesis, however, requires more testing.

The Experimental Approach

In contrast to the lack of tests using the observed economic behavior approach, the EUH has been rigorously tested in an experimental setting, and the experiments and tests have been generally well constructed. A landmark in this class of study was Officer and Halter's work with Australian wool producers. In this study and similar ones, the focus was on the auxiliary hypothesis: can we in fact accurately measure a decision maker's utility function?

The von Neumann-Morgenstern model for obtaining utility functions does so by finding a decision maker's point of indifference between the payoff associated with a gamble and a sure outcome. This method of eliciting preferences has been criticized (Young, et al.) because (a) the act of gambling itself may have

utility or disutility for the decision maker, and (b) the decision maker may have preferences for particular probabilities. Including these two hypotheses results in a model somewhat different than the EUH model. The EUH model proposed that decision makers maximize the expected utility of wealth plus income. Letting U represent the utility function and income plus wealth be y, the EUH recommends:

(4.1)   maximize $E[U(y)]$

where E stands for the expectations operator evaluated over all possible action choices. What Officer and Halter suggest is that a better model might be either

(4.2)   maximize $E\{U[y,f(y)]\}$

(the modified von Neumann-Morgenstern model), or

(4.3)   maximize $E\{U[y,g,f(y)]\}$

(the Ramsey model), where g represents a perceived level of gambling and f(y) represents particular probability levels associated with each income plus wealth level.

Lacking a theory that explicitly incorporates g and f(y) into the decision model, Officer and Halter hold them constant and examine predictions for models (1), (2), and (3). The predictions from models (2) and (3) are then compared with predictions using the model described in equation (1).

To conduct the experiment, the authors needed a choice set separate from the one used to measure preferences. Like Lin, Dean, and Moore they constructed one equal to an EV choice set. The members of the choice set consisted of alternative fodder reserves, a decision problem which was familiar to the decision makers. In this they avoided a criticism of later tests that the choice set was unrealistic and, therefore, not of interest to the participants.

Using carefully measured utility functions described in equations (1), (2), and (3), predictions from thirteen possible fodder reserve choices were obtained. The first test was to compare these predictions with actual choices made. And since the choices were restricted to the EV set, the conditions for the test were not met. The second test was to compare predictions from alternative hypotheses, that decision makers select action choices on the basis of minimum expected cost, model (1), model (2), or model (3). The result was that the Ramsey model, model (3), gave accurate predictions 76% of the time and was superior to both models (1) and (2). The criterion of minimizing expected cost gave accurate predictions only 58% of the time. Moreover, the authors found that after reconsideration and reapplication of models (2) and (3), their accuracy improved. But before reconsideration, expected cost minimization sometimes performed better than model (2), and always better than model (1).

The evidence supports rejection of the naive EUH model. Without attention to decision makers' attitudes toward gambling and probabilities, the model does not predict any better than its competitors such as expected cost minimization models.

There is, however, a disturbing feature of the Officer and Halter study and later ones that use the Ramsey model to predict preferences. If utility of wealth is not independent of probability measures, then applying the EUH would produce unbiased results only if the action choices are described by uniform probability density functions. Applying the EUH over generalized probability density functions would affect both the weighting of the utility as well as utility. This lack of independence between probabilities and income would bias the resulting expected utility measures. Notwithstanding this bias, reasonably accurate predictions were obtained, demonstrating that as a practical tool the EUH cannot be rejected. The Officer and Halter study also suggests that the EUH may often fail

to predict accurately because the auxiliary assumption of accurate utility measurement was not satisfied.

Still the question remains: can the experimental approach based on responses to hypothetical questions be reliably used to obtain utility functions or to test the predictive ability of the EUH in important real world decisions? Binswanger's study provided some answers. He found that before participating in gambles with significant outcomes, decision makers demonstrated different degrees of risk aversion when they played hypothetical outcome games than in an actual game with significant outcomes. Once having participated in the gambling experiment with significant outcomes, though, there was no statistically significant difference between response to hypothetical choices and choices with singificant outcmes.

This result demonstrates that learning does occur in an experiment with significant outcomes. A question which remains is whether or not experience with actual outcomes is required for the learning. Officer and Halter and Webster also observed learning without exposing the respondents to actual outcomes. Probably the most famous learning experience reported in the literature was by Savage, who agreed to alter his action choice when confronted with evidence that he violated the independence axiom of the EUH.

At issue still is whether or not the learned responses to gambles more nearly match actual choices. And, is there a similar learning curve which alters actual responses to economic choices if the choices are made repeatedly? More studies like Binswanger's are required to answer such questions.

Returning to the number of arguments in a utility function question, King and Robison assumed that decision makers maximize an expected utility model with known arguments income + wealth, and unknown arguments, $X_1, \ldots, X_n$, which may or may not be held constant. They argued that decision makers maximize

(4.4) $E\ U(y, X_1, \ldots, X_n)$.

Arguments X_1 through X_n are not measured. The model instead measures

(4.5) $E\ U(y,\varepsilon)$,

where ε is an error term resulting from failure to hold constant or measure variables $X_1, \ldots, X_n$. This approach suggests that decision theorists are naive to believe a single-valued, single-argument utility function can capture all of the information that is needed to predict preferences, or that they can predict a single preferred choice from a choice set with perfect accuracy while accounting for only one argument in the utility function. Using an efficiency criterion developed by Meyer which is consistent with the EUH, King and Robison measured an interval around risk preferences, where risk preferences are measured according to the Pratt-Arrow absolute risk aversion function. Recognizing that their measurements are only accurate in terms of quantifiable probability measures, they offer a somewhat unique approach toward risk measurement. First, they identify as a Type I error the rejection of the preferred choice from a choice set and as a Type II error the failure to order pair-wise comparisons of action choices. Since the EUH with a single valued utility function discriminates on the basis of any absolute difference, it has the greatest likelihood of committing a Type I error and a small chance of committing a Type II error. That is, given any choice set it will select only one choice; the probability that this may not be the preferred choice is the likelihood of a Type I error. On the other hand, all pair-wise choices will be ordered so that the probability of a Type II error is nearly zero. Efficiency criteria such as first-degree and second-degree stochastic dominance have a lower likelihood of producing a Type I error but may result in large Type II errors because of their failure to distinguish preferences.[4]

---

[4]First degree stochastic dominance orders action choices into efficient and inefficient sets for decision makers who prefer more to less. Second degree stochastic dominance orders action choices into efficient and inefficient sets for risk averse decision makers.

As an alternative, the authors proposed an interval that allows a tradeoff between Type I and Type II errors. The interval that King and Robison measured can be of any width or shape. The larger the width, the larger will be the Type II error and the smaller the Type I error. While the interval procedure may avoid many of the problems of discovering a utility measure consistent with the preference orderings of individuals, its ability to do so effectively is related to the width of the interval. Furthermore, methods for determining the optimal width are not fully developed at this time. A major benefit of the interval approach is that it is much easier to apply since it only requires an ordering of action choices, not the discovery of indifference points.

To test Giere's first condition using the interval approach, three questionnaires were administered to a group of graduate students in agricultural economics. The first questionnaire measured risk intervals of different widths at different income levels. The second questionnaire derived utility functions using the modified von Neumann-Morgenstern model with neutral probabilities. The third questionnaire presented decision makers with a series of choices between pairs of distributions. The experiment required that the risk interval measures and the utility functions predict actual choices in each case; the test was the comparison of the predicted choices with the actual.

In this study, the EUH model predicted correct choices 65% of the time, or a 35% Type I error. This evidence alone is a marginal pass of Giere's condition one; it also ordered choices 100% of the time for a zero Type II error. The largest interval predicted correct choices (i.e., did not reject the preferred choice) 98% of the time, while the smallest interval predicted correct choices 75% of the time, or a 2% and a 25% Type I error, respectively. The largest interval meanwhile ordered choices 9% of the time, and the smallest interval ordered them 91% of the time for Type II errors of 91% and 9%, respectively.

The conclusion from the King and Robison study is that the EUH is a useful, but not a perfect, predictive tool. Also, their study points to an important question: what is the optimal tradeoff between Type I and Type II errors and what factors affect this optimal tradeoff? More discriminating models will likely come at a cost of increased Type I errors. Less discriminating models will increase Type II errors.

Experimental Approach with Significant Outcomes

Two studies have been reported recently using the experimental method with significant outcomes approach (Binswanger; Grisley and Kellogg). Both studies were similar in that they constructed an artificial choice set using significant outcomes and measured risk aversion using a partial risk aversion measure (Zeckhauser and Keeler; Menezes and Hanson). Both studies found a distribution of risk aversion measures. While much useful information was obtained in these studies, it is disappointing that no evident test of the EUH was produced. To have conducted a test of Giere's condition one would have required an additional choice set be constructed, different from the first. Could they, for example, have derived utility functions using significant outcomes and then used these to predict actual economic choice? Perhaps a future study could follow such an approach.

Examining the Axioms

A different approach towards testing the EUH is to examine the axioms that define rational behavior and ask whether they conform to observed behavior. The answer that can be readily given is: no, they do not conform, at least not all the time. Consider some of the evidence.

Kahneman and Tversky, summarizing years of research, report consistent violations of the axioms underlying the EUH. In its place they propose a new theory -- prospect theory. The first violation of the rational behavior axioms

is what they refer to as the certainty effect, the overweighting of outcomes that are considered certain. The effect was first demonstrated by the French economist Allais (Allais and Hagen).

In this result, the authors only confirm what Officer and Halter, Webster, and Haneman and Farnsworth found -- that probabilities deserve a place in the utility function. However, rather than attempting to hold the influence of probabilities constant, these authors propose an explicit form for its inclusion. Their proposed model is multiplicative; probabilities are weighted by a function v and outcomes by a utility function U. The resulting ordering index model can be written as,

(4.6)   maximize $E\{U(y)v[f(y)]\}$.

Rather than proposing methods to measure the new function v, they suggest instead that it is a standard function across individuals even though it is not well behaved at its end points. While this new model is intriguing, it lacks two things: a method for measuring the function v and an experiment for testing conditions one and two. Nevertheless, its ability to explain what have been aberrations of the EUH is encouraging.

Machina has also dealt with the Allais and related problems of EUH consistency. In contrast to the prospect theory approach, Machina deduced a version of the EUH without the independence axiom. However, it leaves the EUH as it is now applied as only a local measure using a well-behaved function. Nevertheless, it does resolve many of the inconsistencies easily produced using the EUH.

Another assault on the axioms of rationality underlying the EUH was made by Janis and Mann. They quote John F. Kennedy who asked: "How could I have been so stupid?" after realizing how badly he had miscalculated when he approved the Bay of Pigs invasion (Janis and Mann, p. 657). The EUH portrays a carefully calculated expected utility maximizer weighing each possible alternative. And this

is, in fact, how many decisions are made. But Janis and Mann argue that many other decisions are simply not made in this manner.

They describe five different coping models that describe decision makers' behavior depending on the stress level. (a) Unconflicted adherence: in this model the risks associated with maintaining the status quo are small and, as a result, the status quo is maintained. Subsistence farmers with well-established farming plans may exemplify such a decision model. There is no careful weighing of alternatives, only continued adherence. (b) Unconflicted change: in this model the risk associated with not changing is high, while the stress associated with changing is low. Perhaps this decribes an environment in which past practices have failed. In this model, the action choice selected is the one most salient or the one most highly recommended. Again, there is not weighing of the alternatives, only unconflicted change. (c) Defensive avoidance: the model is characterized by high levels of stress. The decision maker's approach is to shift responsibility, procrastinate, and to remain selectively inattentive to correct information. Because the decision maker does not believe a better solution is available, he fails to examine completely the available alternatives. (d) Hypervigilance (panic): again characterized by high stress levels, the decision maker seizes on hastily contrived solutions, overlooking the full set of consequences because of his excitement. Again, the decision maker fails to act like the EUH rational man. (e) Vigilance: this model is characterized by moderate stress levels. The decision maker carefully assimilates and weighs the information and appraises each choice before making a decision. Only in this model would we expect to find operating our EUH rational man.

Janis and Mann offer no evidence their models meet Giere's condition one. Instead, they emphasize that while not all decision makers reflect the vigilance approach, they would be better off if they did. And the authors offer

suggestions for improving the likelihood that the vigilance (EUH-like) approach will be used. The result is that one more argument probably needs to be added to our utility function -- namely stress. A reasonable hypothesis, but one still in need of testing, is that high or low levels of stress may produce decision making behavior quite different from the rational EUH decision maker.

Preferences for Income and Risk Aversion

In Chapter 2 we introduced the distinction between attitudes towards chance taking and preference for income. We can obtain measures of the latter by arbitrarily indexing a range of incomes and searching for an indifference income. (See equation (3) of Chapter 2). But how does one measure attitudes towards chance taking.

If attitude towards chance taking is a binary trait, then the Ramsey method for utility function estimation accurately measures the utility function--a compound of the preference for income and attitudes towards chance taking. But if the attitude towards chance taking cannot be captured with a binary variable, but in fact depends on the distribution of outcomes associated with the chance, the Ramsey method would not, in any consistent manner, capture ones attitudes towards chance taking ?

Krzysztofowicz has tried to separate the influences of attitudes towards chance taking and preference for riskless income y. We follow Krzysztofowicz and refer to measures of the latter as value functions, $v(y)$, defined over income y. The mapping $w(v(y))$, which accounts for attitudes towards chance taking as well as preference for income which we refer to as the utility function $u(y)$ Krzysztofowicz writes as:

$$(4.7) \quad u(y) = w(v(y))$$

where w is an individuals attitude towards chance taking.

Krzysztofowicz has hypothesized that the transformation y takes the form

$$(4.8a) \quad u(y) = \frac{1-e^{-bv(y)}}{1-e^{-b}} \quad \text{iff } u(y) \quad b > 0$$

$$(4.8b) \quad u(y) = \frac{e^{-bv(y)}}{e^{-b}-1} \quad \text{iff } u(y) \quad b < 0$$

$$(4.8c) \quad u(y) = v(y) \quad \text{iff } u(y) \quad b = 0$$

which hold whenever the decision makers attitude towards risk (b) is constant. The decision maker is therefore constantly risk averse (b>0), constantly risk seeking (b<0) or constantly risk neutral (b=0). Under such circumstances, u is related to v by a unique transformation.

In a series of experiments Krzysztofowicz demonstrates that relative risk attitudes are constant for an individual for a given situation, but neither the value function nor the relative risk attitude is constant across individuals or for the same individual across situations. In very few circumstances is the utility function equal to the value function. These results are distinctly opposed to Kahneman and Tversky's arguments that the value function is constant across individuals and that in many cases u(y) = v(y).

Although Krzysztofowicz's experimental results are valuable in their support of this hypothesis, we are uncomfortable with his assertion that the transformation w is limited to a single functional form each for risk averting, risk preferring and risk neutral individuals. Research to find alternative measures of attitudes towards chance taking and find the form of the function y(.) will continue.

Conclusions

The evidence presented in this chapter could be used to infer that there has been inadequate evaluation of the EUH. The question of interest, can the EUH predict real world decisions, has not been answered satisfactorily because of our own inability to construct a choice set which describes the actual choices facing

the decision maker. Moreover, the prospects are not attractive for constructing a legitimate test of the EUH which would predict action choices from an actual choice set facing the decision maker. Without this actual choice set available to use in a test, we are forced to apply our predictive tests to other experimentally obtained choices. In this type of choice environment the EUH makes correct predictions of experimentally derived pair-wise choices in roughly 60-70% of the cases. We should not expect it to be a perfect predictor because neither the utility function nor the experimentally produced probability density functions used to describe action choices are without error of measurement.

Binswanger and Grisley and Kellogg have come closest to evaluating the EUH's predictive ability in real world settings by constructing choice sets with significant incomes. Unfortunately, since they assumed a utility function rather than derived it, their test really examined whether or not the decision maker was consistent in the manner in which he made choices rather than whether or not the decision maker was an expected utility maximizer. Nevertheless, these studies provided valuable insights regarding the decision process and more studies developed using their methods should be encouraged--and extended to include the explicit measurement of utility functions.

CHAPTER 5

COMPARING INDIVIDUALS' ATTITUDES TOWARD RISK

Introduction

In Chapter 3 we discussed how action choices under conditions of certainty led to unanimity of selection. As long as more is preferred to less, the action choice producing the most is preferred. This is the conclusion of static economics. When uncertainty enters, the unanimity of preference for a particular action choice is lost. The EUH model selected the preferred action choice by maximizing a personalized preference function $U(y)$ over the set of possible action choices.

Since $U(y)$ may differ between individuals, it is necessary to make comparisons between individuals' utility functions in order to discuss differences in preferred action choices. In the examination of preferences reflected by personalized utility functions, we may also want to look for similarities between preference functions of individuals within groups. For example, do decision makers become more adverse to uncertainty as they become older, more wealthy, or more educated?

Obviously, any comparison of the attitudes toward risk of decision makers must begin with the only measure of preferences available -- the utility function $U(y)$. So, we begin by comparing individuals using their personalized utility function $U_j(y)$ j=1, . . . , N where N is the number of individuals being compared. We can then review the literature which has described attitudes toward risk of individuals and in some cases has related them to personal and business characteristics.

Risk Attitudes Inferred from the Shape of U(y)

To compare attitudes towards risk, the standard approach is to ask how different individuals would respond when faced with identical action choices with risky outcomes. Suppose, for example, that a lottery with outcomes y_L and y_H was being offered for sale and that we were in a position to observe the maximum bids of N individuals. The maximum bids represent certainty equivalents which the decision makers would willingly exchange for the lottery. So, at an indifference point, the utility of the certainty equivalent y_{CE}, (a maximum bid price) is equal in preference to the expected utility of the lottery.

For the i-th individual this equality could be written as:

$$(5.1) \quad U_i(y_{CE}) = pU_i(y_L) + (1-p) U_i(y_H)$$

where p is the likelihood with which the decision maker perceives that y_L will occur. We represent this indifference graphically by drawing an arbitrary function $U_i(y)$ in Figure 5.1. The mean of the lottery is $\bar{y} = py_L + (1-p)y_H$. The linear function is expected utility for all possible values of $0 < p < 1$. For p equal to zero, $EU_i(y)$ is $U_i(y_H)$. For p equal to one, $EU_i(y)$ is $U_i(y_L)$. For p such that $\bar{y}$ is the mean, the expected utility is $EU_i(y)$ which is equal to $U_i(y_{CE})$.

The concavity of the function $U_i(y)$ suggests that the average or the expected value of the lottery must exceed its purchase price. The difference between the expected value of the lottery and the certainty equivalent of the lottery is often referred to as π, the "risk premium" (Pratt). It is also customary to order individuals according to their risk premiums -- the larger the risk premium the more risk averse the individual. This ordering procedure, however, has limitations which will be discussed.

The utility function $U_i(y)$ drawn in Figure 5.1 is an arbitrary one. We discussed in Chapter 3 how any linear transformation of $U_i(y)$ would have yielded
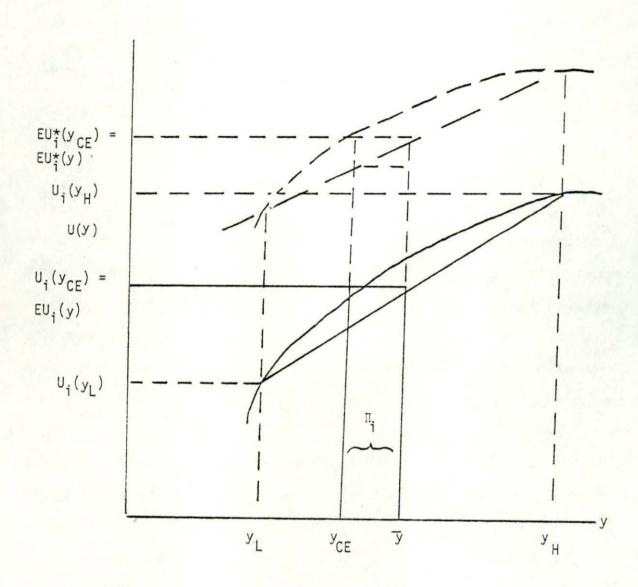
FIGURE 5.1

AN ARBITRARY UTILITY FUNCTION $U_i(y)$ SHOWING INDIFFERENCE
BETWEEN y_{CE} RECEIVED WITH CERTAINTY AND THE LOTTERY WITH
OUTCOMES y_L AND y_H OCCURRING WITH PROBABILITIES
p AND $(1-p)$ RESPECTIVELY

the same indifference. For example, a horizontal shift in $U_i(y)$ to $U_i^*(y)$ also produces indifference between $U_i^*(y_{CE})$ and $EU_i^*(y)$ as the dotted lines in Figure 5.1 illustrate. What produces differences in the risk premium is the bending of the function $U(y)$. In Figure 5.2 we compare two individuals with utility functions $U_i(y)$ and $U_j(y)$. As they are drawn, $U_i(y)$ bends at a greater rate than does $U_j(y)$. As a result, the risk premium π_i associated with $U_i(y)$ is larger than π_j which is associated with function $U_j(y)$. This result might lead us to infer that individual i is more risk averse than individual j.

The utility functions in Figure 5.2 are bending downward. As the function bends less in a downward or negative direction, the size of the risk premium decreases -- in Figure 5.2 the decrease is from π_i to π_j. As the rate of bending in a negative direction approaches zero, the function $U(y)$ approaches a straight line and the risk premium π approaches zero. Thus the certainty equivalent of a decision maker with a linear utility function (with a positive slope) is the mean of the lottery. Because this individual requires no risk premium, he is referred to as risk neutral.

Positive bending of the function $U(y)$ produces, as we would expect, negative risk premiums, or amounts in excess of the mean which decision makers willingly pay to acquire lotteries. We refer to individuals with negative risk premiums as risk preferrers or risk lovers.

The direction of the bending -- negative, zero, or positive -- is indicated by the second derivative of $U(y)$. For $U''(y) < 0$ the bending is negative, while $U''(y) = 0$ indicates no bending and $U''(y) > 0$ implies positive bending. So either $U''(y)$ or the sign on the risk premiums can be used to classify decision makers into the broad categories of risk averse, risk neutral, or risk loving. But the magnitude of the second derivative cannot be used for interpersonal comparisons of risk aversion because an individual's utility function is only unique up to a
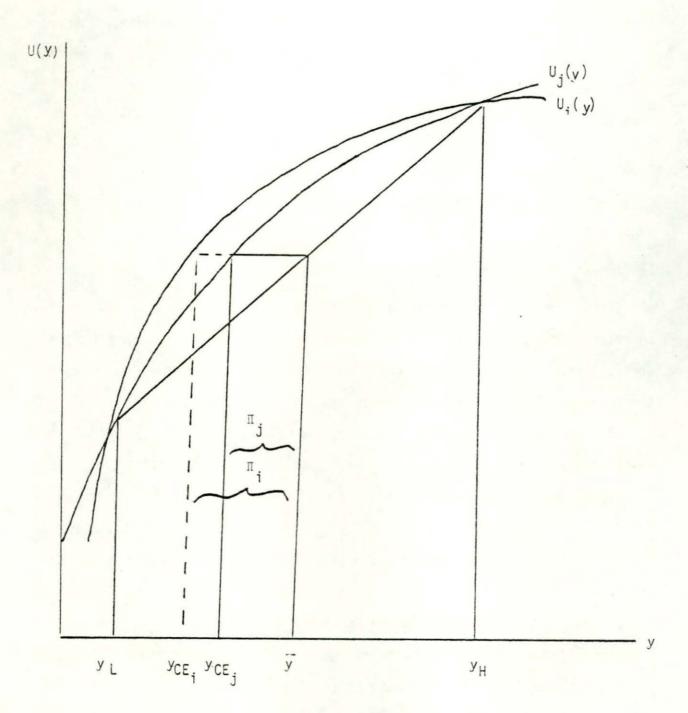
FIGURE 5.2

A COMPARISON OF RISK ATTITUDES OF INDIVIDUALS i AND j WITH
UTILITY FUNCTIONS $U_i(y)$ and $U_j(y)$ AND CERTAINTY
EQUIVALENT INCOMES y_{CE_i} AND y_{CE_j} RESPECTIVELY

positive linear transformation. Thus, the value of the second derivative can be arbitrarily varied by multiplying the utility function by a positive number.

A measure which is unique, which measures not only the direction of bending of $U(y)$ but also the rate of change in slope, is the absolute risk aversion coefficient. Introduced independently by Pratt and Arrow, it is defined as:

$$(5.2) \quad R_a(y) = \frac{-U''(y)}{U'(y)} .$$

A related measure, the relative risk aversion coefficient, $R_r(y)$, measures the elasticity of marginal utility and is defined as:

$$(5.3) \quad R_r(y) = \frac{- U''(y)}{U'(y)} y .$$

Both measures are unaffected by arbitrary transformations of the utility function. They are positive for risk averse decision makers, zero for risk neutral decision makers and negative for risk loving decision makers. Moreover, their uniqueness permits interpersonal comparisons.

The absolute risk aversion function, the measure most often used, like the function $U(y)$ has y as its argument. As a result, for every function $U(y)$ there is a corresponding function $R(y)$. All linear transformations of $U(y)$ would, of course, map into the same function $R(y)$. Thus, there is some advantage to representing decision makers by the magnitude of their absolute risk aversion function $R(y)$ rather than a nonunique utility function $U(y)$.

Comparisons of Risk Attitudes in the Small and in the Large

So far we have inferred that the essence of a decision maker's attitude towards risk is captured by the rate of bending in the ordinal utility function $U(y)$ or the absolute risk aversion function $R(y)$. This function alone, however, has no element of uncertainty or risk included in it. $R(y)$, for example, is simply a function defined over y. But the manner of its derivation through finding indifference between risky alternatives, makes it unclear whether the

function represents simply ordinal ranking of certain income or whether it is also a measure of attitudes toward risk taking. Whatever the truth of the matter -- it is what we use to compare attitudes towards risk of individual decision makers. Moreover, to understand risk attitude comparisons, we must understand how the measure $U(y)$ or $R(y)$ are being used.

The first important distinction is the one made by Pratt between risk attitude measures in the small versus the large. Since $R(y)$ is a function, it is defined over all y and risk attitude measures could be made at any particular point on the function y. Let us choose some specific value for y, call it $\bar{y}$, and ask: "For the i-th and j-th individual, who is more risk averse at income $\bar{y}$?" Another way to ask the question is to ask: "For small gambles with variance σ^2 and mean y, which individual would pay the larger risk premium π to eliminate uncertainty?"

To answer the question just posed, Pratt derived the approximate relationship below:

(5.4) $\pi = R(y)\sigma^2/2$.

Interpreted, the equation reads -- the risk premium π is equal to the value of the absolute risk aversion at $\bar{y}$, the mean of the gamble, times the variance of the action choice divided by 2. The certainty equivalent, or the certain income which provides the same satisfaction as the gamble, can be found by replacing π with $\bar{y} - y_{CE}$ expressed as:

(5.5) $y_{CE} = \bar{y} - R(y)\,\sigma^2/2$.
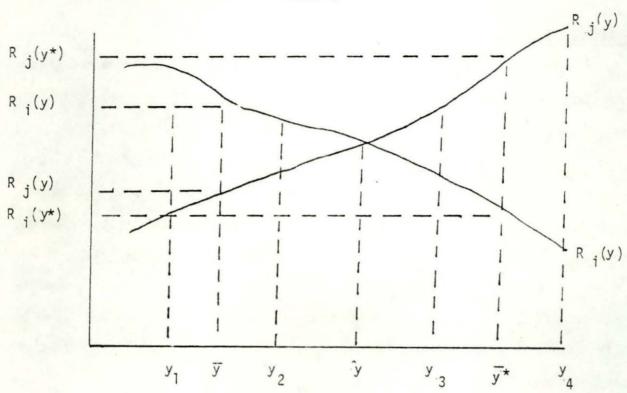
From equation (5.5) we can infer that the more risk averse an individual, a trait indicated by the size of the absolute risk averse function $R(y)$, the larger is the risk premium they would require. So, in the small, or at a point, individuals can be ordered according to the degree of risk aversion by either their absolute risk aversion function valued at a point or by the size of the insurance premium.

Ordering individuals in the large, however, creates another problem. When can we say one individual is always more risk averse than another? For example, consider two individuals i and j whose absolute risk aversion functions $R_i(y)$ and $R_j(y)$ are described in panel A of Figure 5.3. Also assume they are facing an action choice with possible outcomes y_1 and y_2 whose mean outcome is $\bar{y}$. From Pratt's approximation formula we can determine that the i-th individual is more risk averse since $R_i(y)$ is larger than $R_j(y)$. If, however, the action choice has outcomes y_3 and y_4 with mean outcome $\bar{y}*$ then the j-th individual is more risk averse since $R_j(\bar{y}*)$ is greater than $R_i(\bar{y}*)$.

Now suppose the i-th and j-th individuals face a lottery consisting of y_2 and y_3 with mean y. Which one is more risk averse? We cannot say, based on the local or small measure of risk aversion. The individuals could be interrogated to find their respective certainty equivalents, and thus obtain risk premiums for the action choice with outcomes y_2 and y_3. However, we cannot infer that the individual with the larger risk premium is the more risk averse, because many utility functions with corresponding absolute risk aversion functions may have identical risk premiums. In our example, by shifting the probability weights between outcomes y_2 and y_3 we can reverse the risk averse orderings of the i-th and j-th individuals. This is inconsistent with the notion that the risk aversion attitudes are independent of probability measures.

If $R_i(y)$ were greater than $R_j(y)$ as shown in panel b of Figure 5.3, then the risk premium for individual i would always exceed that of the j-th individual, no matter what the probability distribution of action choices. In this case, the i-th individual is globally more risk averse than the j-th individual. In other words, the i-th individual is everywhere more risk averse.

Panel a

Panel b
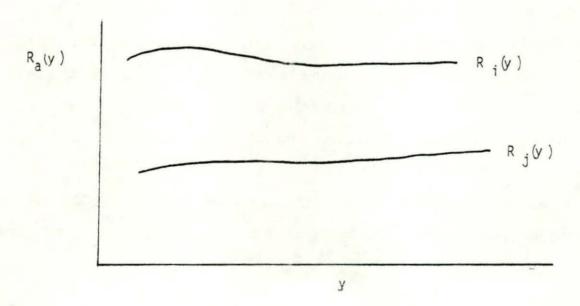
FIGURE 5.3

A COMPARISON OF ABSOLUTE RISK AVERSION FUNCTIONS $R_{ai}(y)$ and $R_{aj}(y)$
OVER OUTCOMES y FOR THE i-th AND j-th INDIVIDUALS RESPECTIVELY

An Alternative Risk Attitude Measure

An alternative measure of risk attitude is obtained by comparing the trade-off between expected values and variances of action choices selected by individuals from an expected value-variance (EV) efficient set. Describing the choice set in terms of means and variances has been a popular approach for several reasons. First, quadratic programming methods can generate the set. And secondly, for risk averse utility maximizing decision makers facing distributions that are normal the preferred choice will always be a member of the EV set.

Figure 5.4 illustrates an EV set. The solid line ACB represents the efficient set. Dots below ACB represent other feasible choices, each described by their respective expected values and variances which are all less preferred than some point on line ACB for risk averse individuals.[1]

To say that C with expected value $\bar{y}_c$ and σ_c^2 is the action choice most preferred by an individual is to argue that it maximizes the individual's expected utility. Let the expected utility for the individual at that point be k. Then we might map an iso-expected utility line equal to k which identifies action choices described in terms of their expected values and variances and represents them as the dotted line tangent to C in Figure 5.4. It is possible, of course, that other iso-expected utility functions might also maximize expected utility at point C. In particular, a straight line could be drawn tangent to C. This dotted line is represented as DCE.

It has been a common practice (e.g., Binswanger; Brink and McCarl; Dillon and Scandizzo) to infer risk attitude orderings based on the slope or trade-off

---

[1]For a more complete discussion on this point, the reader is referred to Markowitz's pioneering article, and more recent articles by Tsiang and Tobin.

between the expected value and variances at the equilibrium action choice. Using such an ordering scheme, individuals who selected action choices above C would be considered less risk averse than those selecting action choices below C. However, such an ordering may not make clear the distinction between risk aversion measures in the small and risk aversion measures in the large.

Consider, for example, the equation for the tangent line at C in Figure 5.4. At the intercept D, an action choice with zero variance has a certainty equivalent outcome y_{CE}. Meanwhile, the slope is a constant times the variance. If we define the constant slope coefficient as $\lambda/2$ we have an equation of the form

(5.6) $\quad \bar{y} = y_{CE} + \lambda\sigma^2/2$

where the intercept y_{CE} plus the slope times the variance at equilibrium equals the expected value of the action choice at point C. We can rearrange the equation to obtain

(5.7) $\quad \bar{y} - y_{CE} = (\lambda/2)\sigma^2$.

Since $\bar{y} - y_{CE}$ is by definition the risk premium Π, we are left with Pratt's local approximation formula given in equation (5.4). The slope is merely the local absolute risk aversion function value at $\bar{y}$. And this being a local risk aversion measure, we may not be justified in making global inferences about risk attitude differences based on EV slope coefficients. Only in the case where all decision makers have constant absolute risk aversion functions could we make such global inferences about risk attitudes.[2]

Classifying Individuals According to Their Risk Attitudes

Having introduced the subject of how risk attitudes can be measured in the context of EUH, we are prepared to summarize studies which have measured risk

_____

[2]For constant absolute risk aversion, the straight line tangent given in equation (5.6) over normal distributions is the iso-expected utility line. It follows then, that the larger the equilibrium slope, the greater the risk aversion of the decision maker. (See Freund and also Hildreth for a more formal proof.)
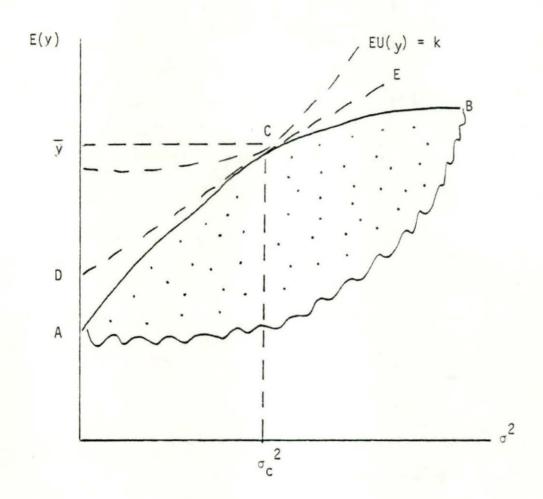
FIGURE 5.4

AN EXPECTED VALUE-VARIANCE EFFICIENT CHOICE SET

attitudes. A comprehensive review was completed by a Western Regional Risk Committee W-149. The committee consisting of D. Young, W. Lin, R. Pope, L. Robison, and R. Selly prepared the material used to construct Table 5.1. The preponderance of the evidence shows that decision makers are risk averse over significant ranges of outcomes. However, where measurement techniques permit- ted, it was not uncommon to observe decision makers displaying risk preferring behavior or a mixture of risk averting and risk preferring behavior.

Several difficulties may be encountered in making a classification such as the one appearing in Table 5.1. For example, Binswanger's and Halter and Mason's studies measured risk attitudes locally. Thus, decision makers must be either risk averse, risk neutral, or risk loving. Mixed behavior would be impossible to observe. Risk attitude may also be influenced by the choice of functions fitted to utility data points. A quadratic utility function must either exhibit risk averting or risk loving behavior. A cubic function will imply mixed behavior. Brink and McCarl, and Moscardi and de Janvry assumed a constant absolute risk averse function which also ruled out mixed behavior.

An alternative approach would be to measure the function $R(y)$ directly without being restricted by available functional forms of $U(y)$ which restrict the function $R(y)$. Such an approach has been implemented by King and Robison. What they estimate is a confidence interval around the decision maker's $R(y)$ function. By selecting local measures of $R(y)$, they constructed global risk attitudes of decision makers.

Summary Measures of Risk Attitudes

Several articles have appeared in our literature over the past several years which have presented summary measures of risk attitudes within an expected util- ity framework. Examples include work by Brink and McCarl, and Bond and Wonder, Dillon and Scandizzo, and Binswanger. In some cases these summary risk measures

TABLE 5.1

DESCRIPTION OF EMPIRICALLY MEASURED RISK PREFERENCES OF
INDIVIDUAL FARMERS FROM THE LITERATURE

	Source	Description of Sample	Sample Size	Percent Distribution of Sample by Risk Classification			
				Averse	Neutral	Preferring	Mixed[a]
1.	Binswanger	Indian farmers and landless laborers	119	71	0	19[b]	--[c]
			117	84	0	9[b]	--[c]
			118	89	0	2[b]	--[c]
			118	97	0	1[b]	--[c]
2.	Conklin, Baquet, and Halter	Oregon orchardists (U.S.A.)	8	37	0	13	50
3.	Dillon and Scandizzo	Brazilian small farmers and sharecroppers	56	70	9	21	--[c]
			47	58	8	34	--[c]
			56	87	0	13	--[c]
			47	79	0	21	--[c]
4.	Francisco and Anderson	Australian pastoralists	21	0	0	5	95
5.	Halter and Mason	Oregon grass seed growers (U.S.A.)	44	33[d]	33[d]	33[d]	--[c]
.	Lin, Dean, and Moore	Large scale California farmers (U.S.A.)	6	50	33	0	17
7.	McCarthy and Anderson	Australian beef ranchers	17	48	29	23	0
8.	Officer and Halter	Australian wool producers	5	60	20	20	0
			5	40	40	0	20
			5	20	0	60	20
			5	80	0	20	0
9.	Webster and Kennedy	Australian sheep and grain farmers	5	80	0	0	20
			5	100	0	0	0
10.	Brink and McCarl	Cornbelt farmers (U.S.A.)	38	66	34	0	--[e]
11.	Moscardi and de Janvry	Mexican peasant farmers	45	100	0	0	--[e]

[a]The risk classification "mixed" includes that portion of the sample having utility functions with both risk averse and risk preferring regions within the relevant range.

[b]Percentages do not sum to 100 because from 2.5 to 10.1 percent of the respondents were classified as "inefficient."

[c]Risk preference classifications were evaluated at a particular point so "mixed" classifications are impossible.

[d]Halter and Mason did not present an exact tabulation of risk preference classifications, but reported "that the number falling into each category was about equal."

[e]"Mixed" classifications were impossible because a constant risk aversion coefficient was assumed by the methodology.
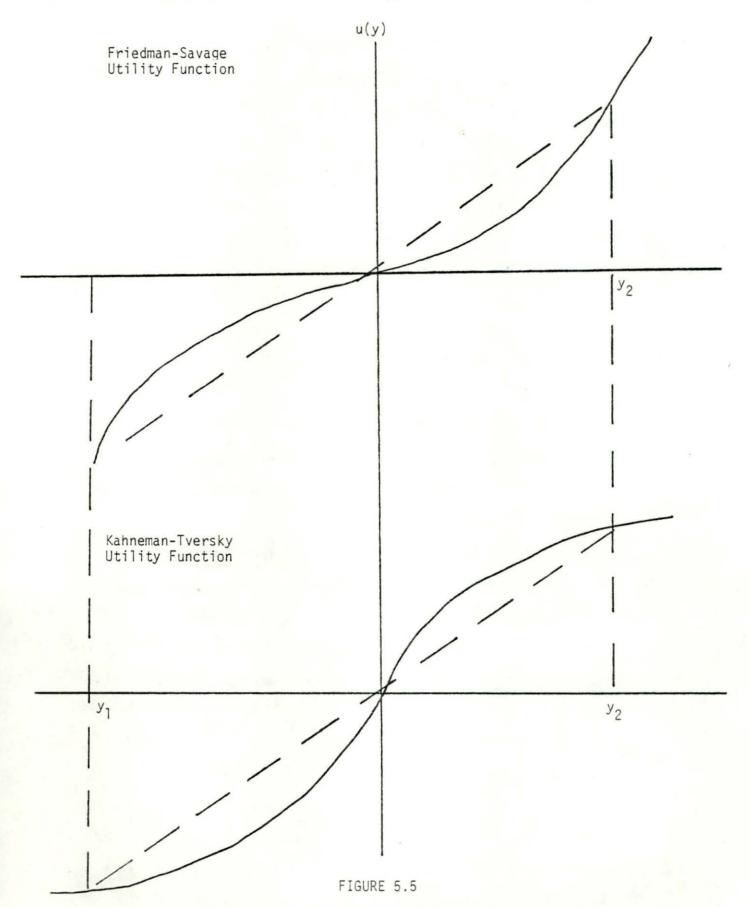
have been used to correlate risk attitudes with other socioeconomic variables in a search for links between the environment and risk attitudes.

Consider, for the moment, the difficulty of obtaining and interpreting summary measures of risk attitudes. To begin, we recognize that all of the risk attitude information available to us is incapsulated in the absolute risk aversion function $R(y)$ which is obtained from the utility function. Measures other than $R(y)$ which may be used to describe risk attitudes have two problems. First the probability density functions of outcomes may be compounded in the risk measure. The second problem is that any summary measure is just that, an incomplete description of the measure. Just as the mean is an incomplete description of the underlying distribution, so will any summary measure of risk attitudes be an incomplete description of $R(y)$.

Still we recognize the need for summary measures noted above using risk premiums, the difference between certainty equivalents and the expected value of outcomes.

Suppose we offered to sell to a set of individuals the same action choice with the likelihood of events being described by a probability density function. We normally infer that the highest bidder for the action choice is the least risk averse since he requires the smallest risk premium. But what if one decision maker's utility function were the backward "S" shaped utility function described by Friedman and Savage while the second decision maker's utility function was the "S" shapled utility function proposed by Kahneman and Tversky.

For the Friedman-Savage decision maker depicted in Figure 5.5 the risky prospect of y_1 and y_2 are equal in utility in the certain outcome of $\bar{y}$. The same is true for the prospect theory decision maker depicted by Kahneman and Tversky. In both cases the risk premium is zero. But in neither case would we say that the

FIGURE 5.5

FRIEDMAN-SAVAGE AND KAHNEMAN-TVERSKY UTILITY FUNCTIONS

decision makers are risk neutral or that they have identical risk preferences. The difficulty is caused by inferring from a single parameter, the risk premium, to a measure in the large, namely the decision makers' utility function.

Summary

In this chapter we have reviewed measures intended to summarize attitudes towards chance taking and preference for income. Assuming the utility function is a compound of the measures and attitudes towards chance taking is a binary variable, the utility function can be uniquely described by Pratt and Arrow's absolute risk aversion function.

Ordering individuals according to the aversion to risk, defined as rate of bending of their utility function, in the large requires one individual's risk aversion function to never lie below the other, a condition which will not permit the ordering of significant number of individuals into risk classes.

As a result, efforts which have categorized individuals according to their risk attitudes have developed alternative summary measures of risk. Sometimes these are consistent with a local measure of the Pratt function of absolute risk aversion.

In some cases these measures have been used to correlate with business or personal characteristics of the decision maker. But the fact that the utility function is a composite of attitudes towards chance taking and preference for income implies little can be learned from such correlations about attitudes towards chance taking.

CHAPTER 6

ORDERING RISKY ACTION CHOICES

Introduction

Uncertain action choices were defined in Chapter 2 as ones whose outcomes are not definitely known. Moreover, we argued that it was the decision maker's knowledge base which determined whether or not action choices were uncertain or certain--and thus we argued that discussions of uncertainty are always subjective. Risky action choices on the other hand were defined as uncertain ones where the outcomes could alter the well being of the decision maker. And since well being is a subjective sensation interpreted by an individual, so is the perception of riskiness.[1]

Action choices may be either uncertain or certain depending on the decision makers sureness of the outcome. Thus we do not talk about action choices as being more or less certain. But riskiness of action choices is another matter. Individuals can and do distinguish between action choices based on their perception of differences in riskiness. And this perception of riskiness is what determines maker's preference for one outcome over another.

Thus for one decision maker to assert that one action choice is riskier than a second reflects his preference for the probabilistic distribution of the outcomes of the first relative to the second. But based on this single decision maker's preferences, we could not infer that individuals in general would prefer

---

[1]Most often riskiness is interpreted in a negative sense; namely, the possibility of material or social loss or injury. This connotation, however, is much too strict for our purposes since its limits potentially uncertain outcomes to ones whose entire range of outcomes reduce the well being of the decision maker. Our definition of risky action choices includes the possibility that outcomes may either improve or reduce the decision maker's well being.

one distribution to another unless they all had their risk attitudes in common. Thus when we talk about riskiness of action choices, we must first talk about attitudes towards risk or well being held by decision makers.

In Chapter 5, we discussed how preferences, measures of well being or risk attitudes reflected by individual utility functions could be used to compare individuals and to order them according to aversion to risk. In this chapter we explore in greater detail the link between risk attitudes and characteristics of probability density functions for action choices which can be used to order action choices according to riskiness. But the basic principle remains. Riskiness of action choices cannot be inferred without specifying the attitudes of the decision maker(s) for whom the ordering is being made. Therefore any riskiness comparisons between risky action choices must be preceded by a statement describing the preferences of the relevant class of decision makers.

Ordering Action Choices According to Riskiness

Having discussed in the previous chapter how individuals can be classified according to measures obtained from their utility functions, we are prepared to examine what implications these utility function measures have on the ordering of action choices. In exploring how orderings of action choices correspond to risk attitudes we will rediscover a familiar relationship: the more complete the ordering, the more demanding will be the preference information. However, information about preferences is not measured without error, so the more complete the ordering, the greater will the chance of incorrect orderings. If we can be satisfied with a less than complete ordering which divides action choices into efficient and inefficient sets, our preference information requirement will be reduced as well as the chance of our obtaining a large Type I error. If, in pairwise comparisons of action choices, fewer orderings are made, there will be a reduction in the error of claiming one distribution is preferred to another when

that statement isn't true. The cost, unfortunately, will be that fewer distributions will be ordered; i.e., more often the test result will be that no preference of one distribution for the other can be determined. Failure to order when an ordering would have been made by the decision maker is referred to as a Type II error.

Utility Function Ordering Rules

A decision maker's utility function contains all risk preference information which is available. Thus this information is the basis for any complete orderings of action choices. Let $U(y)$ be the utility function of a decision maker and let $f_1(y)$, $f_2(y)$, ..., $f_n(y)$ be probability distributions describing the likelihood of outcomes for n risky action choices facing the decision maker. The decision maker's problem is to order them according to their riskiness or synonomously, to order them according to his preferences. This he does by forming the preference indices: $EU(y_1)$, $EU(y_2)$,..., $EU(y_n)$ where y_1, y_2, ..., y_n are the random variables associated with distributions $f_1(y)$, $f_2(y)$,..., $f_n(y)$. These indices then form the basis of a complete ordering of the action choices since any absolute difference in the index can be used to order.

When faced with the problem of ordering action choices according to the preferences of an identified decision maker, an explicit answer is required and the complete ordering performed by the utility function is warranted. However, such an explicit ordering only applies to a single decision maker. Beyond that it has no application.

There have been those, however, who have argued that individuals have similar utility functions described by a common function. Daniel Bernoulli was probably the first to make such a claim, arguing that preferences were described by the logarithmic function:

(6.1) $U(y) = \log y$

This led to the result that distributions with the highest geometric mean were most preferred and the remaining distributions could be ordered according to the relative magnitudes of their geometric means.[2]

The utility function which Bernoulli's log function was to replace was the linear one described earlier as:

(6.2) $U(y) = ky$

where y is a positive constant. Such a rule we have already determined leads to the preference index:

(6.3) $EU(y) = k\bar{y}$

---

[2]The ordering equivalence of the geometric mean criterion and the expected log utility function requires that each criterion yield the same orderings of action choices. Thus if there exists a positive monotonic transformation equating two functions, the orderings will be identical.

The geometric mean $\bar{y}_g$ of outcomes y_1, ..., y_n with likelihoods of occurrences p_1, ..., p_n respectively is:

$$\bar{y}_g = \prod_{i=1}^{n} y_i^{P_i}$$

A log function is a monotonic transformation which applied to the above expression yields:

$$\sum_{i=1}^{n} P_i \log y_1 = \log \left(\prod_{i=1}^{n} y_i^{P_i} \right)$$

Since the expected value of the log utility function is a monotonic transformation of the geometric mean, it must provide the same ordering.

where $\bar{y}$ is the expected value of y. Then if faced with the problem of ordering the n action choices the resulting preference indices would be:

$$k\bar{y}_1, \ k\bar{y}_2, \ \ldots, \ k\bar{y}_n$$

where $\bar{y}_1, \bar{y}_2, \ldots, \bar{y}_n$ correspond to the expected values of the probability distributions $f_1(y), f_2(y), \ldots, f_n(y)$. These preference indices then would be ordered according to their expected values and independent of the positive values of k.

The opposition to such explicit ordering rules is the lack of evidence to support the claim that these functions represent, in general, preferences of individuals.

As a result, efforts have been made to specify decision maker's utility functions more generally. If decision makers were risk averse, then a quadratic utility function of the form:[3]

$$(6.4) \qquad U(y) = y + by^2$$

could be assumed where b < 0. One might argue for a quadratic function as being a reasonable approximation of any concave utility function.[4]

Taking the expectation of the quadratic utility functions leads us to probability density function characteristics which can be used to order action choices according to riskiness. If y is stochastic with expected value $\bar{y}$ and variance of σ^2, equation (6.4) can be written as:

---

[3] Since utility functions are unique up to linear transformations, we can always transform quadratic functions of the form $U(y) = d + ey + fy^2$, where d, e, f are parameters to obtain the expression above which has the single parameter b.

[4] A second order Taylor series approximation would, e.g., lead to a quadratic function approximation in a neighborhood.

(6.5) $E(y + by^2) = E(y) + bEy^2$

Recalling that σ^2 equals $Ey^2 - (Ey)^2$ we can add and subtract $(Ey)^2$ without altering the equality and obtain:

(6.6) $E(y + by^2) = Ey + b\ Ey^2 - (Ey)^2 + (Ey)^2 = \bar{y} + b(\sigma^2 + \bar{y}^2)$

The above criterion with $b < 0$ implies that the riskiness of action choices is dependent on expected values and variances of action choices. For $b < 0$, an increase in σ^2 holding $\bar{y}^2$ constant increases the riskiness of action choices and reduces their preference for risk averse decision makers. Thus for action choices of equal means and different variances, the action choice with the smallest variance is preferred by all risk averse decision makers or decision makers with diminishing marginal utility.

One might place additional restrictions on the quadratic function by limiting the value of "b." Elton and Gruber, for example, have suggested at least one. Since $U(y) = y + by^2$ describes a quadratic, at the y value such that $U'(y) = 1 - 2by = 0$ the marginal utility becomes negative. Since $U'(y) < 0$ isn't an accepted feature of most preference functions, Elton and Gruber argue that the y value at which marginal utility occurs should be at least some specified distance from the mean value of y. If the minimum value selected is, say 2 standard deviations from the mean $y = \bar{y} + 2\sigma$ and $\bar{y}$ and σ are known, b can be solved for. With an explicit minimum value for b, a more refined criterion than EV can be deduced.

This criterion, referred to as the expected value-variance (EV) criterion is described graphically in Figure 6.1.

Each dot in Figure 6.1 describes the expected value and variance corresponding to an action choice. Take, for example, the action choices A_i and A_j identified in Figure 6.1. A_i and A_j have identified expected values by A_j has a larger variance and is less preferred by risk averse decision makers than A_i. Here
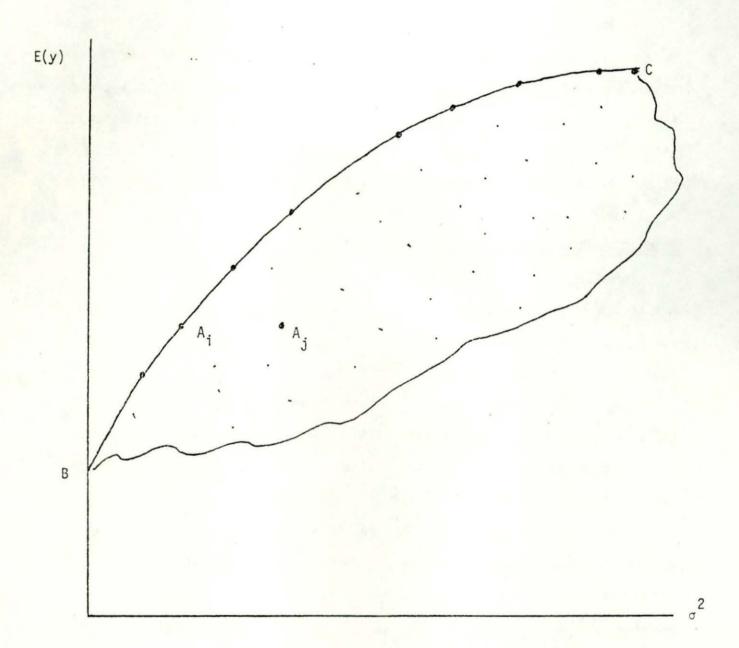
FIGURE 6.1

AN EXPECTED VALUE-VARIANCE EFFICIENT CHOICE SET
REPRESENTED BY THE LINE BC

again, riskier is applied to individuals whose preferences are represented by a quadratic, concave to the origin, utility function.

Action choices along the line BC are preferred to action choices interior to BC for the identified class of decision makers. When action choices are separated into efficient choices (points along BC) and inefficient choices (points interior to BC), there is for each inefficient choice <u>an</u> action choice all decision makers (as defined) would prefer to the inefficient one. To illustrate, all quadratic risk averse decision makers prefer A_i to A_j--therefore A_j is inefficient.

Stochastic Dominance Rules

For some, any functional restriction on the shape of the utility function may be too limiting. As a result, still more general ways to describe decision makers' attitudes towards chance taking have been introduced. The result has been even more generally applicable efficiency criteria.

Without attempting to provide a rigorous proof of the stochastic dominance criteria, an intuitive introduction will be made. For a rigorous discussion of first and second degree stochastic dominance the reader is encouraged to read Hadar and Russell or Hanoch and Levy.

Consider the class of decision makers who prefer more to less, a quite unrestrictive assumption and that a decision maker from this class is faced with action choices A_1 and A_2 whose likelihood of occurrence is described in Table 6.1 below.

The outcomes associated with the action choices are y_1, ..., y_n. The likelihood of the outcomes occurring for action choices A_1 is described by either the density function $f(y)$ or the cumulative function $F(y)$. The likelihood of occurrences of outcomes associated with action choice A_2 is described by either the density function $g(y)$ or the cumulative density function $G(y)$.

Suppose the probability functions f and g are related in the following ways. Outcomes under each have equal likelihoods of occurring except for the ith and kth outcome. Outcome y_k is more likely to occur if action choice A_1 is made while outcome y_1 is less likely. The difference between f and g is that the probability $\alpha > g(y_i)$ has been subtracted from the likelihood of occurrence for the ith outcome under A_2 and added to the likelihood of the kth occurrence.

So for the action choice A_1, an event more satisfying, y_k is more likely to occur at the expense of a less favorable event, y_i, becoming less likely to occur. The result, for all those who prefer more to less, is to make action choice A_1 less risky and more preferred than A_2.

The effect of probability shifts which make an action choice more preferred is demonstrated in the last two columns of Table 6.1. The probability of getting an outcome y or something less (worse) is always less for the action choice A_1 than for A_2. So in general our criterion, called the first degree stochastic dominance criterion (FSD), can be written as follows. The action choice associated with F(y) is always preferred to the action choice associated with G(y) by all decision makers who prefer more to less if the condition

$$F(y) \leq G(y)$$

for all y with strict equality for at least one y. This condition is described graphically in Figure 6.2.

Second Degree Stochastic Dominance

First degree stochastic dominance, which orders action choices into efficient and inefficient sets in accordance with preferences of all decision makers with $U'(y) > 0$, is the most general of the efficiency criterion. The disadvantage, of course, is that for a large number of pairwise action choice comparisons, no preference can be inferred--because all decision makers who prefer more to less must have unanimous preference of one action choice over the other for an

TABLE 6.1

A TABULAR PRESENTATION OF THE LIKELIHOOD OF OUTCOMES
ASSOCIATED WITH ACTION CHOICES A_1 AND A_2 WITH DENSITY
FUNCTIONS f AND g AND CUMULATIVE FUNCTIONS F AND G RESPECTIVELY

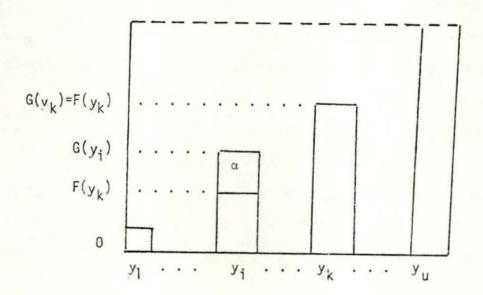| Outcomes | Density Functions | | Cumulative Density Functions | |
	A_1	A_2	A_1	A_2
y_1	$f(y_1)=g(y_1)$	$g(y_1)$	$F(y) =$	$G(y_1)$
.	. .	.	.	.
.	. .	.	.	.
.	. .	.	.	.
y_i	$f(y_i)=g(y_i)-\alpha$	$g(y_i)$	$F(y_i)<$	$G(y_i)$
.	. .	.	.	.
.	. .	.	.	.
.	. .	.	.	.
y_k	$f(y_k)=g(y_k)+\alpha$	$g(y_k)$	$F(y_k) =$	$G(y_k)$
.	. .	.	.	.
.	. .	.	.	.
.	. .	.	.	.
y_n	$f(y_n)=g(y_n)$	$g(y_n)$	$F(y_n) =$	$G(y_n)$

FIGURE 6.2

A GRAPHICAL PRESENTATION OF THE CONDITION ON
CUMULATIVE DISTRIBUTIONS $F(y)$ AND $G(y)$ for $G(y)$
TO BE RISKER OR LESS PREFERRED THAN $F(y)$

ordering to occur. As a result, the number of action choices in a first degree stochastic dominance efficient set is almost always large. And if large numbers of action choices are being generated, e.g., using Monte Carlo procedures, then the criterion becomes unworkable.

The solution is to refine the description of decision makers preferences. In addition to $U'(y) > 0$, we assume risk aversion or $U''(y) < 0$. Defining preferences in this manner leads to a new efficiency criterion called second degree stochastic dominance (SSD).

Again we introduce SSD with an intuitive argument rather than a formal proof. Those desiring a more rigorous presentation are referred to Fishburn, Hadar and Russell, and Hanoch and Levy.

We begin with a comparison between action choices A_1 and A_2 described with probability density function f and g respectively with outcomes $y_1, \ldots, y_n$ arranged in ascending order. The distributions are constructed in such a way that distribution f is obtained from distribution g by shifting probability from the tails to the center of the distribution. For example, $\alpha < g(y_i)$ is shifted from the ith to the (i+1)st outcome. In contrast, $\beta < g(y_{k+1})$ is shifted back from the (k+1)st outcome to the kth outcome. The results are presented in tabular form in Table 6.2.

In Table 6.2 the first probability shift exceeds the shift . For decision makers with diminishing marginal utility can we unequivocably argue the decision maker has been made better off by such a shift? The answer is yes. The shift of α probability which made the more favorable outcome y_{i+1} more likely while reducing the likelihood of the less favorable outcome y_i definitely increased expected utility. In fact it increased by the amount:

(6.7) $\quad U(y_{i+1}) \; \alpha - U(y_i) \alpha = \alpha \Delta U(y_i).$

TABLE 6.2

A TABULAR PRESENTATION OF PROBABILISTIC OUTCOMES OF ACTION CHOICES A_1 AND A_2 EXPRESSED IN TERMS OF PROBABILITY DENSITY FUNCTIONS, CUMULATIVE DENSITY FUNCTIONS, AND THE SUM OF THE CUMULATIVE FUNCTION DIFFERENCES

| | | | Action Choices | | | | |
| | | | A_1 | A_2 | A_1 | A_2 | $\Sigma(G(y)-F(y))$ Sum of |
Outcomes			Probabilities of Density Functions		Cumulative Density Functions		Cumulative Function Differences
y_1	$f(y_1)$	$=$	$g(y_1)$	$g(y_1)$	$F(y_1)$ $=$	$G(y_1)$	0
y_i	$f(y_i)$	$=$	$g(y_i)-\alpha$	$g(y_i)$	$F(y_i)$ $=$	$G(y_i)$	α
y_{i+1}	$f(y_{i+1})$	$=$	$g(y_{i+1})+\alpha$	$g(y_{i+1})$	$F(y_{i+1})$ $=$	$G(y_{i+1})$	α
y_k	$f(y_k)$	$=$	$g(y_k)+\beta$	$g(y_k)$	$F(y_k)$ $=$	$G(y_k)$	$\alpha-\beta$
y_{k+1}	$f(y_{k+1})$	$=$	$g(y_{k+1})-\beta$	$g(y_{k+1})$	$F(y_{k+1})$ $=$	$G(y_k)$	$\alpha-\beta$
y_n	$f(y_n)$	$=$	$g(y_n)$	$g(y_n)$	$F(y_n)$ $=$	$G(y_n)=1$	$\alpha-\beta$

On the other hand a shift in probability from y_{k+1} to y_k which is less favorable than y_{k+1} reduced the expected utility. It is reduced it by the amount

(6.8) $U(y_{k+1}) - \beta U(y_k) = \beta \Delta U(y_k).$

The difference between the gain of expected utility at y_i and the loss of expected utility at y_k can be written as:

(6.9) $\alpha \Delta U(y_i) - \beta \Delta U(y_k) > 0.$

It is greater than zero because 1) $\alpha > \beta$; and 2) diminishing marginal utility requires the marginal utility at y_k be less than y_i or $\Delta U(y_i) > \Delta U(y_k)$.

Thus probability shifts which preserve the sign of the cumulative difference between $\Sigma(G-F) > 0$ will always imply that F is preferred to G.

The cumulative distributions along with the cumulative sum of the differences between F and G are presented in Figures 6.3 and 6.4 respectively. In Figure 6.3 the cumulative distributions differ by probability amount α at y_i ($F(y_i) < G(y_i)$) and by probability amount β at $y_k (F(y_k) > G(y_k))$.

The cumulative sum of the difference between $F(y)$ and $G(y)$ is graphically described in Figure 6.4. This measure is best thought of as the cumulative value of the area between the two cumulative distributions $F(y)$ and $G(y)$. And since they differ only at points y_i and y_k, this area measure will only have two different values, α and $\alpha - \beta$.

Now consider a special application of the second degree stochastic dominance rule which we have just proven. Suppose the choice is between two action choices A_1 and A_2 whose likelihood of outcomes can be described with normal probability density functions with equal expected values and different variances f and g respectively. Two such probability density functions are drawn in Figure 6.5 with their cumulative distributions drawn in Figure 6.6.
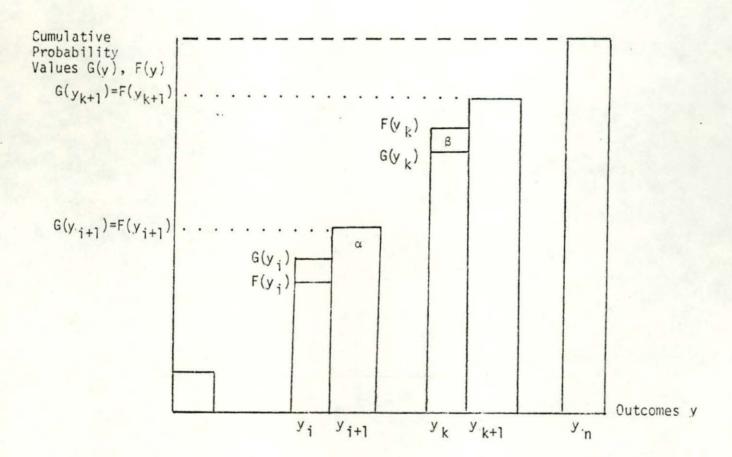
FIGURE 6.3

A GRAPHICAL PRESENTATION OF CUMULATIVE DENSITY FUNCTIONS $F(y)$ $G(y)$ CORRESPONDING TO ACTION CHOICES A_1 AND A_2 RESPECTIVELY
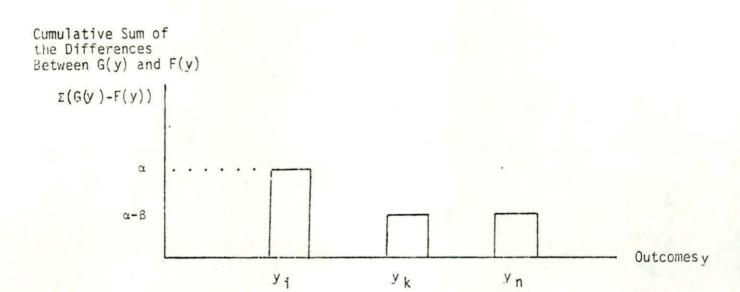
FIGURE 6.4

A GRAPHICAL PRESENTATION CORRESPONDING TO
TABLE 6.2 OF THE CUMULATIVE SUM OF THE DIFFERENCE $G(y)-F(y)$

FIGURE 6.5

PROBABILITY DENSITY FUNCTIONS WITH EXPECTED VALUE
$E_f(y)$ AND STANDARD DEVIATION σ_f COMPARED TO
PROBABILITY DENSITY FUNCTION $g(y)$ WITH
EXPECTED VALUE $E_g(y)$ AND STANDARD DEVIATION σ_g

FIGURE 6.6

CUMULATIVE DENSITY FUNCTIONS F(y) AND G(y)
CORRESPONDING TO ACTION CHOICES A$_1$ AND A$_2$
WHOSE LIKELIHOOD OF OUTCOMES ARE NORMALLY DISTRIBUTED
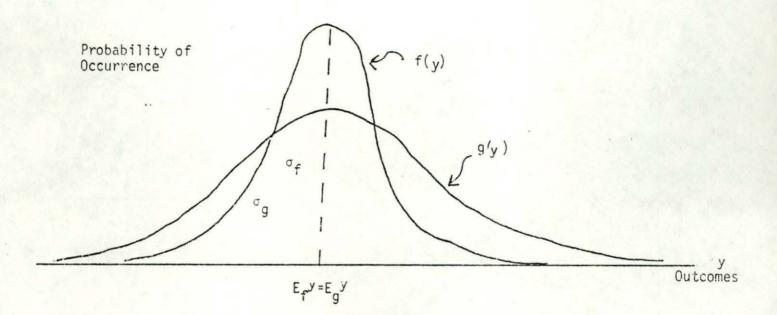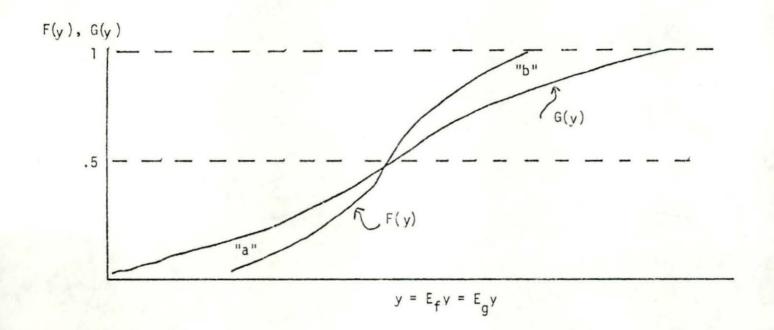
Since the distributions are normally distributed and symmetric, the probability mass is divided equally at the expected values of the distributions. This implies two things: that the cumulative distributions corresponding to F and G will be equal and cross at their expected value outcomes, and that symmetry will require that the difference in area between the distributions to the left of the expected value will equal the difference in area to the right of the expected value. That is area "a" in Figure 6.6 equals area "b." As a result the cumulative sum of the difference between G-F will always be satisfied and F will be preferred to G by all risk averse decision makers. This result leads again to the EV criterion that for distributions of equal means but different variances (standard deviations) the distribution with the smallest variance (standard deviation) is preferred. Thus two justifications for the EV set described graphically in Figure 6.1 are that decision makers have risk averse quadratic utility function or that decision makers are risk averse and face action choices whose outcomes are normally distributed.

Stochastic Dominance with Respect to Functions

The efficiency criteria described above each had the ability to order action choices into efficient and inefficient sets for a particular class of decision makers. To say one set of action choices is preferred to another is to also say it is less risky. Hence the efficiency criteria discussed permitted us to assert some definitiveness to the world risk--or at least how to measure it.

While FSD and SSD add to our understanding of risk and how to measure it, as practical tools they leave much to be desired because of the large size of the efficient set usually obtained. Moreover, the arbitrary classification of decision makers into classes depending on the derivatives of their utility functions may be too restrictive, especially since strong evidence exists that decision makers display both risk loving and risk averse attitudes.

Stochastic dominance with respect to functions is an evaluative criterion which orders action choices for a class of decision makers defined by the analyst rather than being externally imposed. The decision making class is defined in terms of upper and lower bound absolute risk aversion function. In fact, FSD and SSD are special cases of this more general efficiency.

To illustrate, consider the class of decision makers for whom FSD provides an ordering of action choices. They were described as having positive marginal utility $U'(y) > 0$. In terms of the absolute risk aversion function, this placed no bounds since $U''(y)$ was free to take on any value. Thus the decision making class consistent with FSD were $R(y) = -U''(y)/U'(y)$ is defined as:

(6.10) $-\infty < R(y) < \infty$

The SSD set was more discriminating. In addition to $U'(y) > 0$ it required $U''(y) < 0$. Now the function $R(y)$ and the corresponding decision maker class for which it applies is limited to the class of risk averse decision makers for which:

(6.11) $0 < R(y) < \infty$

The lower bound defining the decision making class is the horizontal axis while the upper bound is infinity.

Stochastic dominance with respect to functions allows the decision maker to define functional bounds on $R(y)$; the lower bound function might be $R_1(y)$ while the upper bound function may be $R_2(y)$. Then the class of decision makers is:

(6.12) $R_1(y) \leq R(y) \leq R_2(y)$.

The necessary and sufficient conditions to order action choices consistent with the class restriction above have been proven by Meyer. For distribution F

to be preferred to G by decision makers in the class described by (6.12) the solution procedure requires the identification of a utility function $U_0(y)$ which minimizes

(6.13) $\int_{-\infty}^{\infty} (G(y) - F(y)) U'(y)dy$

    s.t.

    $R_1(y) \leq -U''(y)/U'(y) \leq R_2(y)$ for all y

Equation (6.13) is equivalent to the difference between expected utilities between distributions $G(y)$ and $F(y)$.[5] By minimizing (6.13) a search is made for the decision maker in the defined set least likely to prefer F to G. If a member of the set defined by (6.12) least likely to prefer F to G prefers F to G, then all other included in the set will also and G is eliminated from the efficient set.

If the member of the set described in 6.12 least likely to prefer F to G doesn't, then the procedure is repeated for G relative to F beginning with the search for the individual least likely to prefer G to F.

Meyer's solution to the problem described above is optimal control technique. Details of the solution is given in Meyer and an example is given in King and Robison. Applications of the techniques are illustrated in Love and Robison.

_____

[5]To show this let the difference in expected utility between f and g be:

$\int U(y)f(y) - \int U(y)g(y) = \int U(y)(F(y)-g(y))$

Then applying the change of variable technique to integrate let $dv = f(y)-g(y)$, $U = U(y)$, $U = F(y)-G(y)$ and $du = U'(y)$. Then recalling $udv = uv \mid^{+\infty}_{-\infty} -\int vdu$ we write:

$\int U(y)(f(y)-g(y))dy = U(y)(F(y)-G(y))\mid^{+\infty}_{-\infty} + \int (G(y)-F(y))U'(y)dy$

$= \int (G(y)-F(y))U'(y)dy = \int (G(y)-F(y))U'(y)dy$

Set Stochastic Dominance

One characteristic of all the efficiency criteria presented thus far has been the unanimity of preference requirement. Consider for example the comparison between action choices in A_1 versus A in either Table 6.1 or 6.2. If A_1 dominated A_2 or was preferred, it required that <u>all</u> decision makers prefer more to less in the case of FSD or it required all decision makers with diminishing marginal utility (in the case of SSD) to prefer A_1 to A_2. The same unanimity of preferences is required for stochastic dominance with respect to functions. This unanimity of preference requirement naturally leads to a smaller number of action choices considered to be inefficient than would otherwise be the case if unanimity were not required. Fishburn provided the theoretical breakthrough required to relax the unanimity of preference requirement.

Conclusions

The essence of decision making under uncertainty is to order action choices according to preference. Sometimes this preference ordering is confused with an ordering according to riskiness. But riskiness, we argued earlier, depends on the risk attitudes of the decision maker. Thus any ordering of action choices must of necessity begin by defining for whom a preference ordering is occurring.

A common classification of risk attitudes is to associate them with derivatives of the von Neumann-Morgenstern utility functions. Concave downward utility functions are said to describe risk averse decision makers. Thus orderings which apply generally for risk averse decision makers are said to be orderings according to riskiness as well as preference orderings. Examples of such orderings include second degree stochastic dominance and EV efficiency criterion.

More recent advances in decision theory may force us to rethink the relationship between riskiness and preference. Classes of decision makers can be specified more precisely than by derivative signs of the decision maker's utility function.

The Pratt coefficient provides a natural measure for bounding risk attitudes and defining classes of decision makers. Moreover, the magnitude of this function indicates a willingness to pay for the elimination of uncertainty. When one more risk averese class of decision makers prefers an action choice while another less risk averse group prefers an alternative, we might determine a riskiness ordering of action choices more specific than possible for d viding decision makers according to the sign of $U''(y)$. Unfortunately this ordering depends on the specific definition of the decision maker.

CHAPTER 7

EXTENSIONS OF THE EXPECTED UTILITY HYPOTHESIS

Introduction

Tests of the EUH have focused on its ability to predict farmers preferred action choices. Tversky has argued that in view of the extreme generality of the model on the one hand, and the experimental limitations on the other, the basic question is not whether the model can be accepted or rejected as a whole. Instead, the problem is to discover which of the assumptions of the model hold, or fail to hold, under various experimental conditions.

The three major assumptions of the EUH which concern us are that expected utility maximizers follow the four axioms of rational behavior defined in Chapter 3 (ordering, transitivity, substitution and certainty equivalents among choices), that utilities can be assigned to absolute states of wealth, and that judgments called for in an analysis can be represented accurately by a single, precise number.

Experimental evidence supports the contention that individuals' actions often do not conform with these fundamental assumptions of the EUH. Decision theorists have used this experimental evidence to develop new approaches to understanding decision processes within the general framework of expected utility analysis. Kahneman and Tversky's pioneering work on prospect theory is an attempt to resolve questions arising from the fact that individuals edit information before using it to choose the prospect with the highest value. Because each individual will edit information in unique ways, apparent inconsistencies in preference ordering arise. In addition, Kahneman and Tversky argue that the decision weights which multiply the value of outcomes are determined by factors including, but not limited to, their attendant probabilities.

The independence axiom which underlies the EUH appears to be routinely violated by decision makers. Machina has shown, however, that despite inconsistencies between the independence axiom and actual behavior, the basic concepts, tools and results of expected utility analysis are still applicable. The generalized form of expected utility analysis which he has developed does not require that the independence axiom hold. Instead, all that is required is an assumption of smoothness of preference, and consistency in the shape of utility functions in a given region. An important implication of this weaker assumption is that the shape of the utility function for wealth is a complete characterization of risk aversion whether or not the individual is an expected utility maximizer.

Both of these extensions of the EUH maintain the assumption that individuals can accurately state their preferences in the form of a single number. Proponents of "fuzzy set theory" argue that uncertainty due to randomness and uncertainty due to imprecision and vagueness are both present in decision making. These distinct qualities must be modeled in different ways, the former using probability theory and the latter using fuzzy set theory. Fuzzy set theory provides a means of quantifying the degree of imprecision associated with any input into the decision process through the use of membership functions. The degree of uncertainty or "fuzz" related to an action choice is, therefore, a function of the fuzziness of the inputs.

Prospect Theory

In the remainder of this chapter, the three extensions of the expected utility will be reviewed in more detail beginning with prospect theory. Following Bernoulli, it has generally been assumed that utilities are asigned to states of wealth. Kahneman and Tversky depart from this tradition and analyze choices in terms of changes in wealth rather than states of wealth. They reject the

assumption of classical analysis that preferences reflect a comprehensive view of the options available to the decision maker. Kahneman and Tversky propose instead that people commonly adopt a limited view of the outcomes of decisions: they identify consequences as gains or losses relative to a neutral point. This can lead to inconsistent choices regarding the same objective consequences because they can be evaluated in more than one way depending upon the reference point with which the outcomes are compared.

In developing prospect theory, Kahneman and Tversky cite several violations of the axioms of the EUH. One of these is framing, the effects arising when the same alternatives are evaluated in relation to different points of reference. Framing effects in consumer behavior may be particularly pronounced in situations which have a single dimension of cost and several dimensions of benefit.

In the EUH the utilities of outcomes are weighted by their probabilities. Kahneman and Tversky hold that the decision weights that multiply the value of outcomes do not coincide with the attendant probabilities. Instead, low probabilities are commonly overweighted while intermediate and high probabilities are underweighted relative to certainty. The underweighting of intermediate and high probabilities reduces the attractiveness of possible gains relative to sure ones and reduces the threat of possible losses relative to sure ones. This "certainty effect" leads to violation of the substitution axiom. In prospect theory, an individual's outcome weighting mechanism is represented by a value function. Risk aversion or seeking is explained by the curvature of this function which is usually concave for gains and convex for losses.

The shape of the value function is explained by the "reflection effect" whereby the preferences expressed for negative prospects are the mirror image of those for positive prospects. In other words, the reflecton of prospects around zero reverses the preference ordering. As a result, risk aversion in the

positive domain is accompanied by risk seeking in the negative domain. In conjunction with the certainty effect this leads to risk seeking preference for a loss that is probable over a smaller loss that is certain. This seems to eliminate aversion to variability, at least with respect to losses, as a plausible explanation of behavior. In addition, the function for losses is much steeper than that for gains. If given an equal probability of loosing $y or gaining some amount, individuals usually demand that the potential gains be a multiple of $y before they will engage in the gamble.

To simplify choices, individuals often disregard components that are shared by all prospects under consideration and focus on their differences. This "isolation effect" may produce inconsistent preferences since a pair of prospects can be decomposed in many ways and the different decompositions may lead to different preference orderings.

Prospect theory distinguishes two phases in the choice process. In an initial editing phase, a preliminary analysis of the offered prospects is carried out, often yielding a simpler representation of the prospects. The second phase is one in which the edited prospect with the highest value is chosen. Editing involves several separate actions including coding, where gains and losses are assessed relative to some neutral reference point, combining, where the range of prospects is reduced by combining the probabilities associated with identical outcomes, segregating, where the risky component of a prospect is separated from the riskless component, simplifying, where extremely unlikely outcomes are discarded and other outcomes are rounded, and dominance, where dominated outcomes are rejected.

Many of the apparent inconsistencies in preference ordering result from editing. In the evaluation stage, a decision weight is associated with each probability affecting the impact of probability on the overall value of the

prospect. The resulting value is not a probability measure and the summation of the values is typically less than unity. Using the value function, a weight is assigned to each outcome which reflects the subjective value of that outcome. The resulting set is a measure of the values of deviations from the reference point, or the expected gains or losses associated with each prospect.

Although the evaluation procedure suggested by prospect theory is procedurally similar to that used in expected utility analysis, the two processes are qualitatively different. Prospect theory seeks to explicitly incorporate the subjective impact of probabilities into the utility analysis through the specification of a value function for each individual. The theory also seeks to explain the reasons for apparent inconsistencies found in individual preferences. This descriptive model of preference formation also presents challenges to the theory of rational choice because it is far from clear whether the effects of decision weights, reference points, and framing should be treated as errors or biases, or whether they should be accepted as valid elements of human experience.

Generalized Expected Utility Analysis

Experimental evidence has shown that the independence axiom of the EUH is systematically violated by phenomena such as the St. Petersburg Paradox and the Allias Paradox. Machina argues that, despite these violations, the basic concepts, tools, and results of expected utility analysis are still applicable because they are not dependent upon the independence axiom. They can also be derived from a weaker assumption of smoothness of preferences over alternative probability distributions.

The role of the other axioms of expected utility theory, which amount to the assumptions of completeness and continuity of preferences, is essential to establish the existence of a continuous preference function over probability distributions in much the same way as is done in standard consumer theory. It is the

independence axiom which gives the EUH its empirical content by imposing a restriction on the functional form of the preference function. The independence axiom implies that the preference function may be represented as the expectation with respect to the given distribution of a fixed utility function defined over the set of possible outcomes. In other words, the preference function is constrained to be a linear function over the set of distributions of outcomes, or, as commonly phrased, "linear in the probabilities". For the independence axiom to hold, the local utility functions for all distributions in the range of prospects must be identical. This is often not the case, as will be shown below. This restriction does not apply if we use a generalized form of expected utility analysis proposed by Machina.

Violations of the independence axiom can be demonstrated using the Friedman-Savage utility function. Based on their observations that the willingness of persons of all income levels to buy insurance is extensive and that the willingness of individuals to purchase lottery tickets or engage in similar forms of gambling is also extensive, Friedman and Savage proposed that there is a generalized form of a von Neumann-Morgenstern utility function held by most people (see Figure 7.1). The utility function is concave and implies risk aversion at low income levels, linear and locally risk neutral at the inflection point, and convex and locally risk loving at high income levels. Individuals will be unlikely to take unfair odds in insurance or gambling in amounts close to their initial wealth position given their hypotehsized constant marginal utility for money in this range. Given the chance of significant gains, however, the individual will participate in gambles with unfair odds. The individual will take equally unfair odds for much less in losses than in gains in an attempt to preserve the resources which he holds.
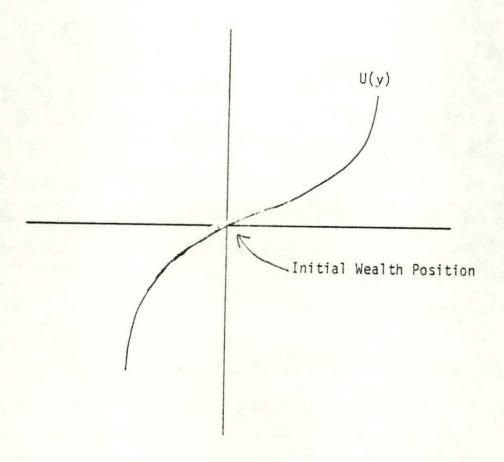
FIGURE 7.1

FRIEDMAN-SAVAGE UTILITY FUNCTION

One implication about human behavior stemming from the assumption of a Friedman-Savage utility function is that people will tend to prefer positively skewed distributions, with larger tails to the right, to distributions which are negatively skewed, with larger tails to the left (Markowitz). There is evidence to suggest that a preference for positive skewness and a relative preference for risk which increases in the upper rather than the lower tails of distributions are also exhibited by global risk averters whose utility functions do not conform to the Friedman-Savage form.

With the later discovery by Markowitz, and Friedman and Savage that the amount an individual would pay for a 1/n chance of winning $ny is an eventually declining function of n, Friedman and Savage modified their utility function to include a terminal concave section. This modified Friedman-Savage utility function is shown below (see Figure 7.2).

Objections were also raised to the original Friedman-Savage form because of the typical response of individuals to a certain type of gamble, known as the St. Petersburg Paradox. The paradox stemmed from the observation that an individual typically would never forego a singificant amount of wealth to engage in a gamble which offered a payoff of $2^i with probability 2^{-i} even though the expected winnings from this gamble are infinite. But the Friedman-Savage function which is consistent with the restrictions of the independence axiom shows, unrealistically, that an individual would take this gamble. The Friedman-Savage form of the utility function is not the only one which suffers from this shortcoming. Menger has shown that whenever the utility function is unbounded, gambles with infinite certainty equivalents can be constructed. Arrow demonstrated that individuals with unbounded utility must violate the continuity or transivity axiom as well as the independence axiom. By bounding the utility function, as is done in the modified Friedman-Savage utility function, the degree of risk aversion is no longer monotonic with respect to outcomes.
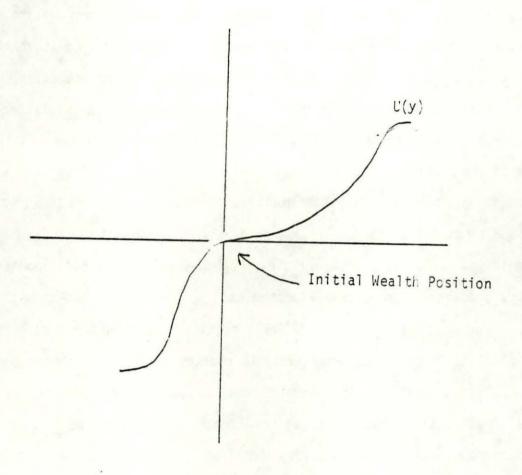
FIGURE 7.2

MODIFIED FRIEDMAN-SAVAGE UTILITY FUNCTION

A third objection to the Friedman-Savage utility function, and one which clearly demonstrates systematic violation of the independence axiom, comes in the form of the Allias Paradox (Allias). The paradox is that individuals systematically rank a stochastically dominating pair of prospects according to a utility function which is more risk averse than the one used to rank a stochastically dominated pair. This is clearly a violation of the independence axiom.

The Allias Paradox can also be used to demonstrate another violation of the independence axiom in that individuals have been found to be oversensitive to changes in the probabilities of low probability, outlying events. This violation has been analyzed by Machina, Kahneman and Tversky, Hagen, and MacCrimmon and Larsson. To compensate for the violation of the independence axiom stemming from oversensitivity to certain probabilities, both psychologists and economists have suggested the use of subjective expected utility models. (See Prospect Theory above.) Although these models allow for a relatively straightforward estimation of the individual's relative sensitivity to changes in low versus high probabilities, Machina argues that they exhibit many undesirable properties. Once the measure of subjective probability is non-linear, behavior is no longer characterized by the shape of the utility function alone and the main results of expected utility theory, such as the characterization of risk aversion by the concavity of the utility function, no longer apply. Subjective expected utility models are also incapable of incorporating the property of monotonicity. This necessarily results in cases where an individual maximizing with a non-linear preference function will prefer some distributions to ones that stochastically dominate them. Similarly, no subjective expected utility maximizer can exhibit general risk aversion even over restricted ranges of possible outcomes (Grether and Plott).

A possible objective to this and other criticisms of EUH models is that when individuals are shown how their choices violate the independence axiom, they then alter their preference so as to conform with it. While this is strong testimony to the normative appeal of the axiom, it is irrelevant to the positive theory of behavior towards risk.

The generalized form of expected utility analysis proposed by Machina does not require that the independence axiom hold. In addition, it leads to results consistent with the Allias Paradox and the St. Petersburg Paradox without requiring the use of subjective probability models. Using local utility functions which display the appropriate qualitative property (e.g., risk aversion) for every local function in a region, the preference function will display the corresponding behavioral property throughout the region. This will occur even if the local utility functions are not the same, or in other words, the individual is not an expected utility maximizer. An important implication of this weaker assumption of smooth preferences is that the concavity of a cardinal function of wealth is a complete characterization of risk aversion in the sense that any risk averter must possess concave local utility functions whether or not he or she is an expected utility maximizer. Thus, the researcher who would like to drop the restrictions of the EUH and study the nature of general risk aversion can apparently still work completely within the framework of expected utility analysis.

Fuzzy Set Theory

Central to the paradigm of decision analysis using the expected utility hypothesis is the often unstated assumption that each of the judgments called for in an analysis can be represented accurately by a single, precise number. Thus, the EUH only addresses uncertainty due to randomness and not uncertainty due to vagueness or imprecision. Much of the unease exhibited by potential users of the

tools of decision analysis stems from concern about their ability to provide sufficiently precise inputs regarding probabilities and utility to receive reliable answers. Watson, Weiss and Donnell argue that probabilities and utility can inherently only be represented by somewhat rough sets of numbers. Their "fuzzy decision analysis" method is motivated by the need to handle the imprecision accompanying the judgmental inputs to decision analysis in a systematic and self-consistent manner.

Zadeh, one of the first to argue for a new fuzzy approach to systems analysis and decision making under uncertainty, holds that imprecision and uncertainty are distinct qualities which must be modeled in different ways, the former using fuzzy set theory and the latter using probability theory. Fuzzy set theory is, therefore, not an alternative to probability theory and the EUH, but a parallel calculus to be used to handle the imprecision inherent in human cognitive processes. The central concept in fuzzy set theory is the membership function which numerically represents the degree to which an element belongs to a set. The function is valued between zero and one and is assessed subjectively with small values representing a low deree of membership in the set and high values representing a high degree of membership. In other words, the statement that "it will probably rain tomorrow", would have a higher degree of membership in a set regarding likelihood of rain than the statement "it might rain tomorrow". Often the values used to represent degrees of membership in a set are not elicited directly. Instead, they are taken from curves drawn by individuals to represent their degrees of belief that an event will occur.

The calculus of fuzzy sets is based on three propositions to which numbers indicating membership should conform. These propositions are analogous to those used in conventional set theory and include:

1. The degree to which X belongs to set A and to set B is equal to the smaller of the individual degrees of membership.

2. The degree to which X belongs to either set A or set B is equal to the larger of the individual degrees of membership.

3. The degree to which X belongs to (not A) is one minus the degree to which X belongs to A.

The calculations involved in the decision analysis can be considered to be a functional relationship between the inputs regarding probabilities and utilities and the output of the analysis in the form of the expected utility of an action. The three relationships cited above are used to deduce the "fuzz" on the output given the fuzziness of the inputs.[1]

As with conventional utility analysis, probability distributions may be generated which characterize the range of possible outcomes for each action choice. Whereas the distributions obtained from conventional analysis are taken to be the true distributions, in fuzzy set theory the extent to which the distribution of inputs, probabilities, and utilities implies an action choice is only as large as the least level of implication for each set. Unless one distribution clearly dominates another, it cannot be said to indicate the preferred action choice. To determine the preferred action choice when two sets overlap, one must determine the extent to which one set is preferred over the other through the use of Zadeh's fuzzy calculus.

There remain questions regarding the axiomatization of a fuzzy set calculus which can be used to elicit membership functions. Experimental evidence does show, however, that individuals are able to draw curves or probability

_____

[1] For particulars of the mathematical methods used, see Watson, Weiss and Donnell and Freeling.

distributions to represent their perceived imprecision regarding degrees of belief such as "better than even", "pretty likely", or "about y%". The precise shapes of these distributions are somewhat arbitrary, but this fact does not affect the inferences which can be drawn from fuzzy set analysis as it is the general shape of the distributions that matter.

The extent to which these extensions to the EUH model will be accepted and adopted remains to be determined. So far, they have not altered in significant ways the "business as usual" of economists. Before this happens several questions must be answered. Can we measure preferences in the context of any of the three extensions listed? Can we build analytical models, which so far have been deduced almost entirely using EUH models? And last, will the extension provide an increase in accuracy commensurate with the cost of more complicated techniques? The answers to these questions will determine the future importance of the extensions to the EUH which this chapter reviewed.

CHAPTER 8

CONCLUSION

In this volume we have explored both the foundations and frontiers of decision theory. This chapter will summarize the strengths of the theory as it stands today and point out weaknesses which provide exciting areas for future research for building a more robust theory of decision making under uncertainty.

The State of the Art

Decision analysis of any type assumes that the decision maker has more than one action choice available to him. Furthermore, it is assumed that the decision problem he faces can be collapsed to comparisons of available actions described in terms of the subjective probability density function of the outcomes associated with each respective action choice. These subjective probability density functions can be described by their expected value, mean, and variance.

Traditionally, we have accepted that if two action choices have the same expected value, the one with the larger variance is more risky. But in order to determine the decision maker's relative preference for the action choices available to him, we must develop some ordering rule. One of the fundamental assertions of this report is that the most suitable ordering rule is one which takes into account the decision makers' attitudes towards risk. To use such an ordering rule we must first determine the individuals' attitudes towards risk.

Many ordering rules have been developed which assume that all decision makers share similar extreme attitudes of optimism or pessimism in response to uncertainty. These include the maximax rule in which the action choice corresponding to the best of the best possible outcomes is preferred and the maximin rule in which the action choice corresponding to the best of the worst possible outcomes is preferred. These are ineffective criterion because they ignore the many other possible outcomes and probabilities with which they may occur.

An alternative to these models is the safety-first model. In its simplest form the safety-first model focuses on a safety or disaster level of outcome such as the income below which a firm will go bankrupt or the minimum crop yield needed to meet subsistence requirements. Whatever the interpretation of the level of safety, this model assumes the objective of selecting the action choice so that chances of experiencing outcomes below the level are minimized.

In contrast to the specific outcome focus of the safety-first model is the expected utility hypothesis (EUH) which allows each outcome which influences the well being of the decision maker to influence the preference index. The EUH asserts that if a decision maker's behavior is consistent with a set of axioms of rational behavior, they will weight outcomes specified in term of income or wealth, y, according to a peronalized utility function U(y). The expected value of this function for each action choice then provides a single value preference ordering index.

The measurement of an individual's preferences requires the assumption that he can identify the most and least favorable outcomes of any action choice. These extreme outcomes are then used to construct a series of gambles over the range. By adjusting either the value of the outcome or its probability of occurrence, a point of indifference between two gambles can be obtained. After a sufficient number of indifference points are obtained, a utility function can be derived using either statistical or graphical methods. The utility function can be used to weight the expected outcomes of each action choice, and the resulting expected utilities are used to determine preference ordering. Individuals will prefer the outcome with the highest expected utility.

One of the reasons we are interested in discovering individuals' utility functions is that this may allow us to rank individuals according to their attitudes towards risk. We may also want to look for similarities in attitudes toward risk within groups which share certain socioeconomic characteristics.

The most commonly used method of comparing individual attitudes toward risk is to rank them according to their response to an identical lottery. We can interpret their indifference point in the gambles as their expected utility of the lottery. The difference between the expected utility of the lottery and its certainty equivalent is often referred to as the risk premium. The more risk averse an individual, the larger his risk premium will be. This provides a basis for ranking individuals.

But, because an individual's utility function is only unique up to a positive linear function, the risk premium approach does not give us a complete characterization of the individual's attitude towards risk even in the region of the gamble. Alternative measures which are unique over the range of the gamble and which incorporate not only the general shape of the utility function but the rate of change of its slope are the Arrow-Pratt coefficients of absolute risk aversion and relative risk aversion. Both measures are unaffected by arbitrary transformations of the utility functions.

The Arrow-Pratt coefficients provide us with the information necessary to rank individuals according to their risk preferences over the range of monetary outcomes covered by the specific gamble. But this does not tell us whether one individual is globally more risk averse than another. The rankings of individuals obtained over different regions of their utility functions can vary widely. For example, in an identical gamble involving $1.50, individual A may be more risk averse than individual B, while in gamble involving $150.00, individual B may be more risk averse than A. Without some rule for ordering individuals over their entire preference functions, we can say little about interpersonal comparisons of attitude towards risk. For the same reason, it is difficult to classify any one individual as risk averse, risk neutral, or risk loving over the entire range of his utility function.

Most studies of decision makers' behavior measure risk attitudes locally. This allows them to be classified as risk loving, neutral or averse. Often the classification is influenced by the choice of functions fitted to the utility points as many functions are restrictive in the type of behavior they can exhibit. Our ability to order action choices into efficient and inefficient sets for a particular class of decision makers has been greatly enhanced by the development of stochastic dominance rules.

Despite the widespread use of the EUH to determine decision makers' preference orderings, many questions regarding its accuracy still remain. Among these concerns are whether decision makers' preferences are actually revealed in a game-like setting, the intertemporal validity of utility functions, and whether a theory which includes income as its only independent argument can be usefully applied in real world situations. Most studies which have used the EUH to predict decision makers' action choices tell us little about its robustness. To meet the requirements of a good test of a theoretical hypothesis one must show that the predicted outcome conforms with the actual outcome and that the same accuracy of prediction could not be attained through an alternative model.

To test the EUH, two choice sets are required: one for use in deriving the individual's utility function, and one to predict his utility maximizing choice. Three approaches have been used to construct the required choice sets: the actual economic behavior approach, the experimental approach, and the experimental approach with significant outcomes. Although the experimental approach with significant outcomes meets many of the objections raised regarding the other two approachs, it is quite expensive. Therefore, the experimental approach which elicits the individual's utility function through a series of hypothetical gambles is most commonly used.

Although these studies provide some support for the EUH, they also lead us to some disturbing conclusions. First, decision analysts may be naive to believe that a single-valued, single-argument utility function can capture all of the information needed to predict choices or that they can predict a single preferred action choice from a choice set. Second, decision makers' attitudes towards gambling and probability may affect their elicited utilities and need to be incorporated into the model. Recent work by mathematical psychologists has shown that utility functions elicited using common methods are actually compound utility functions for individuals' preference for riskless income and attitude toward chance taking. Third, the assumption that probability measures are independent of wealth may be unjustified. Fourth, we have found that choices among artificial lotteries are affected by learning but we have not yet determined whether learning occurs which alters actual responses to economic choices if the choices are made repeatedly. Lastly, considerable doubt has been raised regarding the assumptions that decision makers follow the axioms of rational behavior and that judgments called for in an analysis can be expressed accurately through the use of a single, precise number.

In response to this last concern, decision theorists have used experimental evidence to develop new approaches to decision modeling with the general framework of expected utility analysis. Prospect theory has attempted to resolve questions arising from the fact that individuals do not always follow the axioms of rational behavior assumed by the EUH. The diversion of actual behavior from that assumed by the EUH is due, in part, to the fact that individuals have unique ways of editing information before using it to determine the expected value of the outcome of an action chocie. In addition, prospect theory argues that decision weights which multiply the value of outcomes are comprised of a set of factors which include, but are not limited to, probabilities.

Although the independence axiom of rational behavior is consistently violated by decision makers, development of a generalized form of expected utility analysis still allows us to use the basic concepts, tools, and results of expected utility analysis. This is accomplished through the use of a weaker assumption of smoothness of preferences, or consistency in the shape of utility functions in a given region. An important implication of this weaker assumption is that the shape of the utility function for wealth can be used as a complete characterization of risk attitude whether or not the individual is an expected utility maximizer.

Both prospect theory and the generalized form of expected utility analysis continue the EUH's focus on uncertainty caused by randomness. Fuzzy-set theory has been developed as a parallel calculus which models the uncertainty which is introduced due to imprecision or vagueness in human cognitive processes. Fuzzy-set theory allows for representation and quantification of the degree of imprecision or vagueness associated with any input into the decision process through the use of membership functions.

A Final Note

One may ask why an understanding of risk attitudes and how to measure them is important or, where choices between risky actions are required, why we do not confront decision makers directly. Although most decisions are made by the direct involvement of decision makers in the decision process, in an important number of cases, the direct involvement of decision makers is not possible. Computer simulation models may generate thousands of action choices which a single decision maker could not possibly subjectly evaluate. Having a decision criteria to reduce the choice set presented to decision makers is extremely helpful.

In other cases, policy makers may be called on to make decisions which affect large numbers of individuals. Knowing the risk attitudes of those affected and their likely response can be valuable information in policy design. And if risk attitudes can be found to correlate closely with various socioeconomic variables, policies and decisions affecting groups not directly involved in the decision process may be even more finely tuned.

Finally, we may study decision making under uncertainty because decision makers' decision skills can be improved. Learning, although not discussed in this report, is almost always an important by-product of application of decision making rules.

Problems which will continue to plague decision theorists in the context of the EUH are numerous. Fuzzy set theory, prospect theory, various safety-first models, and Machina's generalized expected utility models are all challengers or some might say extensions to the existing EUH framework. All claim to model well some results which apparently contradict results obtained from the EUH. That there are shortcomings in EUH applications should not, however, come as too great a surprise. No one has claimed to be able to measure with perfect accuracy either decision maker's utility functions or the probability distributions which describe their feasible action choices. All EUH single value index values will, therefore, contain errors. Whether or not these errors in measurement eliminate the EUH as a practical decision tool is still under debate. We tend to agree with Dillon, however, that it remains a practical and useful decision tool in most applied situations. For a review of methods and results of recent applications of decision theory in agricultural settings see Fleisher and Robison, forthcoming.

Still, there is room to improve and some of the newly proposed models may offer improvements. But to replace the EUH will not be easy because it offers

scientists in the decision theory field and economics a unique combination of tools. It is, according to Hey, the basis of at least 95 percent of disciplinary models in risk analysis including the literature applicable to decision makers on small farms. For these models, the precision implied by the EUH in ordering action choices is required. Comparative statistics involving maximization techniques of calculus requires a single valued, precisely measured and described function. Without such functions, the disciplinary progress made thus far would have been impossible. And this tool will not likely be discarded because it doesn't work precisely in all applied situations.

So, we will continue to work with the EUH model. Efficiency criteria which separate action choice sets into one containing the expected utility maximizing choice and one that doesn't allows us to be less demanding in our risk measurement applications. Yet, our theory continues to assume such precise measures exist. The importance of understanding risk and its role in decision making, the solid foundation already built, and the opportunities to make new discoveries combine to make risk analysis an exciting and dynamic portion of the discipline of economics.

BIBLIOGRAPHY

Anderson, J. R., J. L. Dillon and B. E. Hardaker. Agricultural Decision Analysis, Iowa State University Press, Ames Iowa, 1977.

Allais, M. "The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School," Part II of Expected Utility Hypothesis and the Allais Paradox. Dordrecht, Holland, D. Reidel Publishing Co., 1979.

Arrow, K. J. Essays in the Theory of Risk Bearing, Chicago: Markheim Publishng Co., 1971.

_____. "The Use of Unbounded Utility Functions in Expected Utility Maximization: Response," Quarterly Journal of Economics, 88(1974):136-138.

Baumol, W. J. "An Expected Gain Confidence Limit Criterion for Portfolio Selection," Management Science, 10(1963):174-81.

Bell, David and Howard Raiffa. "Marginal Value and Intrinsic Risk Aversion," Harvard Graduate School of Business Administration Working Paper Number HBS 79-65, 1979.

Bernoulli, D. "Specimen Theoriae Novae de Mensura Sortis," Commentarri Academiae Scientiarum Imperialis Petropoliatanae 1738, Tomus V., pp. 175-92. Translated by Louise Sommer as "Expositions of a New Theory on the Measurement of Risk," Econometrica, 22(1954):23-26.

Bernoulli, J. Ars conjectandi, 1713. Translated into German by R. Haussner as "Wahrscheinlichkeitsrechnung," Ostwald's Klassiker der Exakten Wissenschaften, No. 107 and 108. Leipzig: W. Englemann, 1899.

Bierman, L. "Optimal Gambling Systems for Favorable Games," Fourth Berkeley Symposium on Mathematical Statistics and Probability, Berkeley, 1961.

Binswanger, H. P. "Attitudes Toward Risk: Experimental Measures in Rural India," American Journal of Agricultural Economics, 62(1980):395-407.

Bond, Gary and Bernard Wonder. "Risk Attitudes Amongst Australian Farmers," Australian Journal of Agricultural Economics, April (1980):16-33.

Brink, L., and B. McCarl. "The Trade-off Between Expected Returns and Risk Among Corn Belt Farmers," American Journal of Agricultural Economics, 60(1978)259-63.

Conklin, F., A. Baquet and A. Halter. "A Bayesian Simulation Approach for Estimating Value of Information: An Application to Front Forecasting," Oregon State University Agricultural Experiment Station Technical Bulletin No. 136, May 1977.

Debreu, G. Theory of Value: An Axiomatic Analysis of Economic Equilibrium, New York: Wiley, 1959.

de Finetti, B. Theory of Probability: A Critical Introductory Treatment, Vol. 1. Translated by Antonio Machi and Adrian Smith. New York: Wiley (1970), 1974.

Dillon, J. L. "Bernoullian Decision Theory: Outline and Problems," in Risk, Uncertainty and Agricultural Development, Ames Iowa: Iowa State University Press, 1979.

Dillon, J., and P. Scandizzo. "Risk Attitudes of Subsistence Farmers in Northeast Brazil: A Sampling Approach," American Journal of Agricultural Economics, 60(1978):425-35.

Elton, E. J. and M. J. Gruber, "On the Optimality of Some Multiperiod Portfolio Selection Criteria," Journal of Business, April 1974, 237-48.

Feder, G. "Farm Size, Risk Aversion and the Adoption of New Technology Under Uncertainty," Oxford Economic Papers, 32(1980):263-283.

Fishburn, P. C. "Convex Stochastic Dominance with Continuous Ditribution Functions, Journal of Economic Theory 7(1974).

Fleisher, B. "Decision Theory: A Review and Critique of Applications in Agriculture," unpublished Masters Thesis, Department of Agricultural Economics, Michigan State University, 1983.

Fleisher, B. and L. Robison. "Decision Theory and Attitudes Towards Risk: A Review and Critique of Applications in Agriculture," forthcoming.

Francisco, E. and J. Anderson. "Chance and Choice West of the Darling," Australian Journal of Agricultural Economics, 16(1972):82-93.

Freeling, A. N. S. "Fuzzy Sets and Decision Analysis," IEEE Transactions on Systems, Man and Cybernetics, SMC-10(1980):341-354.

Freund, R. J. "Introduction of Risk into a Programming Model," Econometrica, 24(1956):253-63.

Friedman, M. and L. J. Savage. "The Utility Analysis of Choices Involving Risk," Journal of Political Economy, 56(1948):279-304.

Giere, R. N. Understanding Scientific Reasoning, New York: Holt, Rineart and Winston, 1978.

Grether, D. M. and C. R. Plott. "Economic Theory of Choice and the Preference Reversal Phenomenon," American Economic Review, 69(1979):623-638.

Grisley, W. and E. D. Kellogg, "Risk-Taking Preferences of Farmers in Northern Thailand--Measurement and Implications," paper presented at AAEA annual meetings, Urbana, Illinois, 27-30 July, 1980.

Hadar J. and W. R. Russell, "Rules for Ordering Uncertain Prospects," American Economic Review, 59(1969):25-34.

Hagen, O. "Towards a Positive Theory of Preference Under Risk," in Expected Utility Hypothesis and the Allais Paradox, edited by M. Allais and O. Hagar, Dordrecht, Holland: D. Reidel Publishing Co., 1979.

Halter, A. N. and G. W. Dean. Decisions Under Uncertainty, Dallas: South-Western Publishing Co., 1971.

Halter, A. N. and R. Mason, "Utility Measurement for Those Who Need to Know," Western Journal of Agricultural Economics, 3(1978):99-109.

Hanach, G. and H. Levy. "Efficiency Analysis of Choices Involving Risk," Revie of Economic Studies, 36(1969):355-45.

_____. "Efficient Portfolio Section with Quadratic and Cubic Utility," J. of Business, 43(1970):181-89.

Haneman, W. M. and R. L. Farnsworth. "Risk Preferences and Perceptions in the Use of IPM," paper presented at AAEA annual meetings, Urbana, Illinois, 27-30 July, 1980.

Hey, J. D. Uncertainty in Microeconomics, New York: New York University Press, 1979.

Hildreth, C. G. "Point Estimates of Ordinates of Concave Functions," Journal of American Statistical Association, 49(1954):598-619.

Janis, I. L. and L. Mann. "Coping with Decisional Conflict," American Scientist, 64(1976):657-67.

Jeffreys, H. Theory of Probability, Second Edition, Oxford: Clarendon Press, 1948.

Johnson, G. L. "Contributions of Economists to a Rational Decision Making Process in the Field of Agricultural Policy," in Decision Making in Agriculture, proceedings of International Conference of Agricultural Economists. Edited by T. Dames and K. E. Hunt, Oxford Economics, 1977, p. 34.

_____. "Single Enterpreneur Decision Theory," pp. 123-135 in Proceedings: Workshop on Capital Investment Theory, Palo Alto, California: Electric Power Research Institute, 1978.

_____. "Small Farms in a Changing World," pp. 7-28 in W. J. Sheppard, editor, Proceedings of Kansas State University's 1981 Farming Systems Research Symposium "Small Farms in a Changing World: Prospect for the Eighties," Manhattan, Kansas: Kansas State University Farming Systems Research Paper No. 2, 1982.

Kahneman, D. and A. Tversky. "Prospect Theory: An Analysis of Decisions Under Risk," Econometrica, 47(1979):263-91.

Kaminsky, M. "Comment on Unimportance of Risk for Technology Design and Agricultural Development Policy," in Economics and the Design of Small-Farmer Technology, edited by A. Valdes, G. M. Scobie, and J. L. Dillon, Ames Iowa: Iowa State University Press, 1979.

Kataoka, S. "A Stochastic Programming Model," Econometrica, 31(1968):181-96.

King, R. P. and L. J. Robison. "An Interval Approach to the Measurement of Decision Maker Preferences," American Journal of Agricultural Economics, 63(1981):510-20.

Knight, F. H. Risk, Uncertainty and Profit, Boston: Houghton Mifflin, 1933.

Krzyztofowicz, Roman. "Strength of Preference and Risk Attitude in Utility Measurement," Organizational Behavior and Human Preference, 88(1983):88-113.

Lin, W. W. and H. S. Chang. "Specification of Bernoullian Utility Function in Decision Analysis," American Economic Review, 30(1978):30-36.

Lin, W., G. Dean and C. Moore. "An Empirical Test of Utility Versus Profit Maximization in Agricultural Production," American Journal of Agricultural Economics, 56(1974):497-508.

Love, R. O. and L. J. Robison, "An Empirical Analysis of the Intertemporal Stability of Risk Preferences," Department of Agricultural Economics Staff Paper No. 82-24, Michigan State University, 1982.

Luce, R. D. and H. Raiffa, Games and Decisions, New York, 1957.

MacCrimmon, K. R. and S. Larsson. "Utility Theory: Axioms versus 'Paradoxes'," in Expected Utility Hypotheses and the Allais Paradox, edited by M. Allais and O. Hagen, Dordrecht, Holland: D. Reidel Publishing Co., 1979.

Machina, M. J. "Expected Utility Analysis without the Independence Axioms," Econometrica, forthcoming.

_____. "Portfolio Selection," Journal of Finance, 7(1952):77-91.

Markowitz, H. "The Utility of Wealth," Journal of Political Economy, 60(1952):151-58.

Markowitz, H. "Portfolio Selection," J. of Finance, 7(1952):77-91.

Mason, R. and A. N. Halter. "Risk Attitude and Forced Discontinuance of Agricultural Practices," Rural Sociology, 45(1980):435-447.

Masson, R. T. "Utility Functions with Jump Discontinuities: Some Evidence and Implications from Peasant Agriculture," Economic Inquiry, 12(1974):559-66.

Menezes, C. F. and D. L. Hansen. "On the Theory of Risk Aversion," International Economic Review, 62(1972):77-86.

Menger, K. "The Role of Uncertainty in Economics," in Essays in Mathematical Economics in Honor of Oskar Morgenstern, edited by M. Shubik, Princeton University Press, Princeton, 1967.

Meyer, J. "Choice Among Distributions," Journal of Economic Theory, 14(1977):326-36.

Meyer J. and R. Pope. "Unbiased Estimation of Expected Utility and Agricultural Risk Analysis," mimeo, Texas A&M University, 1980.

Miyamoto, John M. "On the Foundations of Multiattribute Psychophysics," Draft, 1983.

Moscardi, E. and A. de Janury. "Attitudes Toward Risk Among Peasants: An Econometric Approach," American Journal of Agricultural Economics, 59(1977):710-716.

Officer, R. and A. Halter. "Utility Analysis in a Practical Setting," American Journal of Agricultural Economics, 50(1968):257-77.

Pratt, J. "Risk Aversion in the Small and in the Large," Econometrica, 32(1964):122-36.

Pyle, D. and S. Turnovsky. "Safety-First and Expected Utility Maximization in Mean-Standard Deviation Portfolio Analysis," Review of Economics and Statistics, 52(1970):1083-86.

Ramsey, F. P. The Foundations of Mathematics, New York: Harcourt Brace, 1931.

Robison, L. J. "An Appraisal of Expected Utility Hypothesis Tests Constructed from Responses to Hypothetical Questions and Experimental Choices," American Journal of Agricultural Economics, 64(1982):367-375.

_____. "An Explanation of Apparently Risk Preferring Behavior by Risk Averse Decision Makers," mimeo, Department of Agricultural Economics, Michigan State University, 1982.

Robison, L. J., P. J. Barry, J. Kliebenstein and G. Patrick. "Risk Attitudes: Concepts and Measurement Approaches," in Risk Management in Agriculture, edited by Peter J. Barry, Unpublished Manuscript, 1982.

Robison, L. J. and B. Fleisher. "Risk: Can We Model What We Can't Define or Measure?" pp. 34-51 in "Risk Management Strategies for Agricultural Production Firms," proceedings of a seminar sponsored by Southern Regional Research Project S-180 in San Antonio, Texas, March 28-30, 1983.

Rothschild, M. and J. E. Stiglitz, "Increasing Risk: 1. A Definition," Journal of Economic Theory, 2(1970):225-43.

Roumasset, J. A. "Unimportance of Risk for Technology Design and Agricultural Development Policy," in Economics and the Design of Small Farmer Technology, edited by Alberto Valdes, Grant M. Scobie, and John L. Dillon, Iowa State University Press, Ames Iowa, 1979.

Roy, A. D. "Safety-First and the Holdings of Assets," Econometrica, 20(1952):431-49.

Samuelson, P. A. "St. Petersburg Paradoxes: Detanged, Dissected, and Historically Described," Journal of Economic Literature, 15(1977):24-55.

Sarin, Rakesh Kumar. "Strength of Preference and Risky Choice," Operations Research, 30(1982):982-997.

Schoemaker, P. J. H. "The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations," Journal of Economic Literature, 20(1982):529-63.

Spence, M. and R. Zeckhauser. "The Effect of Timing of Consumption Decisions and the Resolution of Lotteries on the Choice of Lotteries," Econometrica, 40(1972):401-3.

Stovall, J. "Income Variation and Selection of Enterprises," American Journal of Agricultural Economics, 48(1966):1575-79.

Telser, L. "Safety-First and Hedging," Review of Economic Studies, 23(1955-1956):1-16.

Tobin, J. "Liquidity Preference as Behavior Toward Risk," Review of Economic Studies, 25(1958):65-86.

Tsiang, S. C. "The Rationale of the Mean-Standard Deviation Analysis, Skewness Preference, and the Demand for Money," American Economic Review, 62(1972):354-71.

Tversky, A. "Utility Theory and Additivity Analysis of Risky Choices," Journal of Experimental Psychology, 75(1967):27-36.

Venn, J. The Logic of Chance, Third Edition, London and New York: Macmillan (1866) 1888.

von Neumann J. and O. Morgenstern, Theory of Games and Economic Behavior, Princeton, 1944.

Watson, S. R., J. J. Weiss and M. L. Donnell. "Fuzzy Decision Analysis," IEEE Transactions on Systems, Man and Cybernetics, SMC-9 (1979):1-9.

Webster, J. P. G. "The Analysis of Risky Farm Management Decisions: Advising Farmers About the Use of Pesticides," Journal of Agricultural Economics, 28(1977):243-59.

Webster, J. and J. Kennedy. "Meauring Farmers' Tradeoffs Between Expected Income and Focus Loss Income," American Journal of Agricultural Economics, 57(1975): 97-105.

Whittaker, J. K. and J. R. Winter. "Risk Preferences of Farmers on Empirical Examples, Some Questions and Suggestions for Future Research," Risk Analysis in Agriculture, Research and Educational Developments, pp. 217-24, Department of Agricultural Economics, AE4492, University of Illinois, June 1980.

Young, D., W. Lin, R. Pope, L. Robison and R. Selley. "Risk Preferences of Agricultural Producers: Their Measurement and Use," Proceedings of 1979 Annual Meeting of Western Regional Research Project, W-119, pp. 1-28, Department of Agricultural Economics, AE4478, University of Illinois, July 1979.

Zadeh, L. A. "Fuzzy Sets," Information and Control, 8(1965):338-353.

Zeckhauser, R. and E. Keeler, "Another Type of Risk Aversion," Econometrica, 38(1970):661-65.