

DECISION THEORY NOTES

CONTENTS

Basic Topics in Decision Theory

- Decision Problems

- Decision Makers

- Decision Making

Decision Rules

- The Maximin Value Rule

- The Minimax Regret Rule

- The Maximum Expected Value Rule

Information and Risk

- Conditional Expected Value and Bayes Strategies

- The Economic Theory of Information

- Utility and Risk

Appendix – Axiom System for Utility Theory

References

ELEMENTS OF DECISION THEORY

BASIC TOPICS IN DECISION THEORY

Decision theory pertains to human decision making in a world of incomplete information and incomplete human control over events. Decision theory posits two players: a cognitive human and a randomizing nature. The human, called the decision maker, performs analyses, makes calculations, and cognitively decides upon a course of action in an effort to optimize his or her own welfare. The metaphorical nature is non-cognitive, does not perform analyses or make calculations, and does not choose courses of action in any self-interested way. Rather, nature blithely selects courses of action purely in a probabilistic way.

Decision Problems

The two fundamental concepts of decision theory are states of nature and acts. States of nature are under the control of nature and beyond the control of the decision maker, and are probabilistically selected by nature. Acts are under the control of the decision maker and any one of the available acts can be selected by the decision maker. Further, decision theory presumes that the problem is presented to the decision maker, i.e., that the problem itself, like the states of nature, is beyond the control of the decision maker. A *decision problem* is represented as a pair $\langle S, A \rangle$ composed of a set S of states of nature and a set A of acts. So specified, $\langle S, A \rangle$ is a *decision problem under uncertainty*. If the decision maker has a probability system over the states, then $\langle S, A \rangle$ is a *decision problem under risk*.

Decision Makers

Decision Theory posits that the human decision maker brings to the resolution of a decision problem beliefs and preferences. Specifically, the theory presumes that the decision maker possesses a probability system that captures his or her (partial) beliefs about nature's selection of states, a belief system about the outcomes accruing to the performance of the acts in the various states of nature, and a preference structure over the outcomes. Thus, the *decision maker* is a triple $\langle P, F, U \rangle$ composed of a probability measure P , an outcome mapping F , and a utility function U .

The probability measure P is defined over the set of states of nature, and captures the decision maker's view of the process whereby the states are selected by nature. The outcome mapping F is defined on the Cartesian product of the states of nature and the acts, and presents the outcome resulting from performing each act in each of the states of nature. Thus, the outcome mapping F generates a new set O of outcomes. The utility function is defined over the set O of outcomes and represents the decision maker's preferences over the outcomes.

In what follows we will consider both decision making under uncertainty and decision making under risk. We will treat the former by simply ignoring the decision maker's probability system, and we will treat the latter by incorporating the decision maker's probability system. The focus of this course, however, is on decision making under risk.

Decision Making

Decision making is composed of a two step process. First, the acts in A are ordered, and second the "best" act is selected. Typically, the first step is completed by assigning a number to each act, and then using the complete ordering property of the real numbers to order the acts. If the outcomes are "goods," like gains, income, etc., then the best act is the act with the highest number. Contrariwise, if the outcomes are "bads," like losses, costs, etc., then the best act is the act with the smallest number.

DECISION RULES

Decision rules are composed of two commands. The first command tells the user how to assign numbers to acts; the second command tells the user how to choose among the numbers assigned by the first command. Decision rules are designed to reflect some human attitude toward decision making. The rules considered here reflect two forms of pessimism and general rationality.

In what follows we use a running example. Consider a very simple decision problem where there are two states and two acts, and where the payoffs are given in the following table:

	a_1	a_2
s_1	100	0
s_2	-50	20

The Maximin Value Rule

One of the most popular rules for decision making under uncertainty is based on simple pessimism. The rule, called the maximin value rule, is a rational response in a world where the metaphorical nature is out to get the decision maker.

Presuming the payoffs reflect a "good," the maximin value rule requires the user to first assign to each act the minimum value that act can receive, and second, to choose the act with the maximum assigned value. The foregoing example yields the following:

First step – the numbers assigned to the acts are -50 and 0, respectively, since these are the worst payoffs for each act.

Second step – select act #2 because 0 is the larger of the two assigned numbers.

The maximin value rule reflects a kind of pessimism wherein the decision maker views nature as being "out to get me," and thereby minimizes the worst that nature can inflict. Put differently, the maximin value rule reflects a kind of pessimism wherein the decision maker always "minimizes the downside risk." This kind of pessimism is often observed among people who survived the Great Depression.

The Maximin Regret Rule

Another popular rule for decision making under uncertainty is also based on pessimism. Called the maximin regret rule, this rule attempts to model pessimism with more sophistication than the maximin value rule.

The maximin value rule assigns a value to each act without considering the values accruing to the other acts. The Maximin Regret rule avoids this problem by replacing the payoffs with regret values. The latter, often called opportunity costs, represent regret in the sense of lost opportunities. Specifically, the regret value that replaces a payoff is the distance between the payoff and the best payoff possible in the relevant state of nature.

The regret table for our payoff table

	a_1	a_2
s_1	100	0
s_2	-50	20

is as follows:

	a_1	a_2
s_1	0	100
s_2	70	0

The regret values for the first state of nature (s_1 = your ticket wins) are calculated as follows:

$$0 = 100 - 100$$

$$100 = 100 - 0.$$

The regret values for the second state of nature (s_2 = your ticket does not win) are calculated as follows:

$$70 = 20 - (-50)$$

$$0 = 20 - 20.$$

The Minimax Regret rule requires the user to first assign to each act the maximum regret that act can receive, and second, to choose the act with the minimum assigned regret value. The foregoing example yields the following:

First step – the numbers assigned to the acts are 70 and 100, respectively, since these are the maximum regret values for each act.

Second step – select act #1 because 70 is the smaller of the two assigned numbers.

The Maximum Expected Value Rule

The most popular rule for decision making under risk is the maximum expected value rule. Here the decision maker employs the identity utility function $U(x) = x$, and thereby adopts the attitude of risk neutrality.

The Maximum Expected Value rule represents rationality in the following way. Suppose the decision maker has a coherent and consistent preference ordering over the outcomes, and suppose the decision maker

also has probabilities over the states of nature. If the decision maker assigns to each act the expected value of the act, then the preference ordering over the outcomes will be represented in the natural ordering of the expected values.

Assuming the payoffs are "goods," the Maximum Expected Value Rule requires the user to first assign to each act the expected value accruing to that act, and second, to choose the act with the maximum assigned value. Suppose the probabilities are $P(s_1) = .5$ and $P(s_2) = .5$. Then the foregoing example yields the following:

First step – the numbers assigned to the acts are 25 and 10, respectively, determined as follows:

$$\begin{aligned} E[V(a_1)] &= P(s_1)F(s_1, a_1) + P(s_2)F(s_2, a_1) \\ &= (.5)(100) + (.5)(-50) \\ &= 25 \end{aligned}$$

and

$$\begin{aligned} E[V(a_2)] &= P(s_1)F(s_1, a_2) + P(s_2)F(s_2, a_2) \\ &= (.5)(0) + (.5)(20) \\ &= 10. \end{aligned}$$

Second step – select act a_1 because 25 is the larger of the two assigned numbers.

There seems to be a problem here. For most individuals, act a_2 is superior to act a_1 . How can the most popular decision rule lead to such a wrong choice? The answer is straight forward – the Maximum Expected Value rule does not model the behavior of a person who is averse to risk. To model such behavior we must introduce the notion of a utility function to account for risk attitude. This we will do shortly. In the mean time, however, we will continue with the Maximum Expected Value rule because it simplifies the presentation of the economic theory of information.

INFORMATION AND RISK

The Maximum Expected Value Rule admits of two primary extensions. First, the rule supports an economic theory of information and the formulation of strategies. Second, through the introduction of utility functions, the Maximum Expected Value Rule becomes the Maximum Expected Utility Rule accounts for the decision maker's attitude toward risk. We will treat these in turn.

Conditional Expected Value and Bayes Strategies

The economic theory of information, due primarily to Jacob Marschak (1974), is a decision theory wherein the objects of choice are information systems. An information system for the decision problem $\langle S, A \rangle$ is a pair $\langle Z, P_{z/s} \rangle$ where Z is a set of signals (related to the states in S), and $P_{z/s}$ is a matrix of conditional probabilities on the signals z in Z given states s in S . A probability of the form $P(z/s)$ is called a reliability probability, and the matrix $P_{z/s}$ represents the quality of the information system $\langle Z, P_{z/s} \rangle$.

Reliability probabilities and the decision maker's own prior probabilities determine, via Bayes' Theorem, the matrix of posterior probabilities $P_{s/z}$. Hence, this variation on decision theory is often referred to as Bayesian decision theory. Bayes theorem is as follows:

$$P(s_i / z_k) = \frac{P(s_i)P(z_k / s_i)}{P(z_k)} = \frac{P(s_i)P(z_k / s_i)}{\sum_j P(s_j)P(z_k / s_j)}$$

The first equality is known as the basic form of Bayes' Theorem and is the theorem proved by Bayes (1763); the second equality is known as the extended form of Bayes' Theorem. The latter is the former with the theorem on total probability used to generate $P(z_k)$ in terms of its component parts.

Consider the following example. Suppose we have a problem with two states where the prior probabilities are $P(s_1) = .5$ and $P(s_2) = .5$, and an information system with two signals where the matrix of reliability probabilities is

	P(z/s)	
	z ₁	z ₂
s ₁	.8	.2
s ₂	.4	.6

In words, the information system is 80% reliable in detecting that the first state will occur and 60% reliable in detecting that the second state will occur. The prior and reliability probabilities yield the following posterior probabilities:

	P(s)	P(z/s)		P(s/z)	
		z ₁	z ₂	z ₁	z ₂
s ₁	.5	.8	.2	.4/.6	.1/.4
s ₂	.5	.4	.6	.2/.6	.3/.4

The posterior probabilities are calculated as follows:

$$.4 / .6 = \frac{(.5)(.8)}{(.5)(.8) + (.5)(.4)} = \frac{.4}{.4 + .2},$$

$$.2 / .6 = \frac{(.5)(.4)}{(.5)(.8) + (.5)(.4)} = \frac{.2}{.4 + .2},$$

$$.1 / .4 = \frac{(.5)(.2)}{(.5)(.2) + (.5)(.6)} = \frac{.1}{.1 + .3},$$

and

$$.3 / .4 = \frac{(.5)(.6)}{(.5)(.2) + (.5)(.6)} = \frac{.3}{.1 + .3}.$$

Note that if the decision maker receives the signal z_1 , then the probabilities $.4/.6$ and $.2/.6$ replace the prior probabilities $.5$ and $.5$, respectively, whereas if the decision maker receives the signal z_2 , then the probabilities $.1/.4$ and $.3/.4$ replace the prior probabilities $.5$ and $.5$, respectively. Therefore, to resolve the decision problem with the information system we just resolve the original decision problem twice, once for the world where the decision maker receives signal z_1 and once for the world where the decision maker receives signal z_2 .

To resolve the decision problem with information we follow the same two step procedure we followed earlier. The first step is to calculate the (conditional) expected values for the acts given each of the signals. The second step is to choose, for each signal, the act with the greater expected value. The first step produces the table of conditional expected values $E[V(a)/z]$, as follows:

	P(s)	P(z/s)		P(s/z)	
		z_1	z_2	z_1	z_2
s_1	.5	.8	.2	$.4/.6$	$.1/.4$
s_2	.5	.4	.6	$.2/.6$	$.3/.4$
		$E[V(a)/z]$			
				z_1	z_2
		a_1		$30/.6$	$-5/.4$
		a_2		$4/.6$	$6/.4$

The second step is to choose the optimal act for each signal. If the decision maker receives signal z_1 , then the optimal act is a_1 because $30/.6$ is greater than $4/.6$; and if the decision maker receives signal z_2 , then the optimal act is a_2 because $6/.4$ is greater than $-5/.4$.

The pair of propositions "if z_1 , then a_1 " and "if z_2 , then a_2 ," usually abbreviated $\langle a_1, a_2 \rangle$, is called the *Bayes Strategy*. Such a strategy constitutes a contingency plan since the decision maker know the proper reaction to each signal before any signal is received. The booklet that came with your car includes various Bayes strategies. For example, most such booklets tell you that "if the warning light for oil pressure is on, then stop the car," and "if the warning light for oil pressure is not on, then keep driving." Similarly, the same booklets tell you that that "if the

warning light for coolant is on, then drive the car to the nearest service station," and "if the warning light for coolant is not on, then keep driving."

The Economic Theory of Information

The economic theory of information is decision theory applied to information systems. Note that economic value is attached to an economic system and not to a message. Put differently, you attach economic value to an information source like the Wall Street Journal and not to the messages that appear in any particular issue. You do so because you attach economic value to reliability.

The logic behind the economic theory of information is straightforward. Since the signals z in Z are random variables, with probabilities determined by (the denominator of) Bayes' Theorem, the decision maker can attach a value to a Bayes strategy by calculating the expected value of the strategy. For our running example, the one-trial experiment EXP_1 generated the Bayes strategy $\langle a_1, a_2 \rangle$, so the expected value of EXP_1 is

$$\begin{aligned} E[EXP_1] &= E[V(\langle a_1, a_2 \rangle)] \\ &= P(z_1)E[V(a_1)/z_1] + P(z_2)E[V(a_2)/z_2] \\ &= (.6)(30/.6) + (.4)(6/.4) \\ &= 36. \end{aligned}$$

The foregoing value, usually referred to as the Gross Value of the information system, can be used in a number of ways. First of all, what we have really done is assign a value to a one-trial experiment based on the information system. Thus, we need to compare the foregoing value with the value from no information, i.e., the value of a zero-trial experiment. The latter is simply the expected value of the optimal act in the resolution of the decision problem without the information system. For the running example, this is $\max E[V(a)] = E[V(a_1)] = 25$. Then the net value of a one-trial experiment is

$$E[EXP_1] - E[EXP_0] = 36 - 25 = 11,$$

and we can compare this latter value to the cost of a one-trial experiment to determine if the experiment is cost effective.

Second, since the gross value of no information, i.e., $E[EXP_0]$, is the same for all information systems, we can compare different information systems by simply comparing the gross value attached to each. Thus, just as the expected value function orders the acts in the set A , the gross value function orders the set of all information systems for the decision problem $\langle S, A \rangle$.

Third, we can use the net economic value measure to determine the optimal sample size for the set of many-trial experiments. To do so, we must produce a reliability probability matrix for the relevant number of trials. This is straight forward if the trials are independent; it is less easy if they are not. Using the data from our running example, and presuming independent trials, the reliability matrix for a two-trial experiment is as follows:

	P(z/s)		P(zz/s)			
	z_1	z_2	z_1z_1	z_1z_2	z_2z_1	z_2z_2
s_1	.8	.2	(.8)(.8)	(.8)(.2)	(.2)(.8)	(.2)(.2)
s_2	.4	.6	(.4)(.4)	(.4)(.6)	(.6)(.4)	(.6)(.6)

Clearly, the middle terms in each row of the $P(zz/s)$ matrix match. Therefore, we can save some time and calculation effort by changing the $P(zz/s)$ matrix to the following:

	P(zz/s)		
	exactly two z_1	exactly one z_1	exactly no z_1
s_1	(.8)(.8)	2(.8)(.2)	(.2)(.2)
s_2	(.4)(.4)	2(.4)(.6)	(.6)(.6)

The connection with the binomial distribution should be obvious, as should the extension of the foregoing theory to the case of the n -trial experiment for any integer n . The economic value of two trials is then determined as follows:

	a ₁	a ₂	P(s)	P(#z ₁ /s)			P(s/z)		
				2 z ₁	1 z ₂	1 z ₂	2 z ₁	1 z ₂	1 z ₂
s ₁	100	0	.5	.64	.32	.04	.32/.40	.16/.40	.02/.20
s ₂	-50	20	.5	.16	.48	.36	.08/.40	.24/.40	.18/.20
							E[V(a)/z]		
							2 z ₁	1 z ₂	1 z ₂
a ₁							28/.40	4.0/.40	-7.0/.20
a ₂							1.6/.40	4.8/.40	3.6/.20

The Bayes Strategy is $\langle a_1, a_2, a_2 \rangle$ and the gross value of the two-trial experiment is


$$E[\text{EXP}_2] = .40(28/.40) + .40(4.8/.40) + .20(3.6/.40) = 36.4.$$

Note that $E[\text{EXP}_0] = 25.0$, $GV[\text{EXP}_1] = 36.0$, and $GV[\text{EXP}_2] = 36.4$. This suggests that as the number of trials n increases, $E[\text{EXP}_n]$ increases but at a decreasing rate. This intuition is correct. Thus, the $E[\text{EXP}_n]$ function has a shape similar to that of the logarithm function, except $E[\text{EXP}_n]$ is bounded above.

The optimal sample size is determined by finding the value of n that maximizes the difference $E[\text{EXP}_n] - C[\text{EXP}_n]$. Suppose the cost of the n -trial experiment involves a fixed set-up cost f and a variable cost with a fixed per-trial cost c . Then

$$C[\text{EXP}_n] = cx + f.$$

This is a linear function with intercept f and slope c . Assuming $0 < E[EXP_0] < f$, we have the following diagram:



The upper bound on $E[EXP_n]$ for all n is $E[EXP]$, i.e., the gross value of the infinite trial experiment. We now encounter a counterintuitive result - $E[EXP]$ is finite for all decision problems and decision makers where the payoffs are finite. The value of $E[EXP]$ is easily determined. A perfect predictor is an information system $\langle Z, P_{z/s} \rangle$ where $P_{z/s}$ is the unit matrix with as many rows and columns as there are acts in A . For our running example, we have the following:

	$P(s)$		$P(z/s)$		$P(s/z)$		
	a_1	a_2	z_1	z_2	z_1	z_2	
s_1	100	0	.5	1	0	.5/.5	.0/.5
s_2	-50	20	.5	0	1	.0/.5	.5/.5
					$E[V(a)/z]$		
					z_1		z_2
					a_1	50/.5	-25/.5
					a_2	0/.5	10/.5

The Bayes Strategy is again $\langle a_1, a_2 \rangle$ and the gross value of the perfect information system is

$$E[EXP] = .5(50/.5) + .5(10/.5) = 60.$$

Since this system is perfect, no information system can have a higher gross value. Thus, $E[EXP]$ serves as an upper bound to the gross value measure.

Utility and Risk

The maximum expected value rule is extended to the maximum expected utility rule by employing a non-linear utility function to represent the decision maker's attitude toward risk. The basic attitudes beyond risk neutrality are risk aversion, usually viewed as the keystone of middle class normalcy, and risk preference, usually view as an abnormality (Fishburn and Kochenberger, 1979). As we will see later, the majority attitude toward risk is a hybrid of these basic attitudes.

Utility theory has evolved over the recent centuries and the contemporary form of the theory was developed in response to the work of Daniel Bernoulli (1738). Therein Daniel Bernoulli solved the puzzle, known as the St. Petersburg Paradox, raised by his cousin Nicholas Bernoulli on 1713. The St. Petersburg Paradox is as follows. Consider a simple game played with a fair coin where the coin is flipped until it comes up tails, whereupon the game terminates. The payoff to the player is \$2 if the sequence is HT, \$4 if HHT, \$8 if HHHT, and so forth. Thus, the player receives $\$2^n$ if the sequence is n Hs followed by a T.

The expected monetary value of the game is an infinite sum of 1's and thereby infinite. Therefore, if the value of a game is its expected value, then a rational individual would be willing to pay a very large amount, indeed any finite amount, to play the game. However, Nicholas Bernoulli observed that most individuals place a small finite value on the game. Daniel Bernoulli, in response to his cousin's observation, concluded that the correct value of a game is not the expected monetary value, but rather the expected utility-of-money value, and that the utility of wealth function is not linear. More specifically, by "arguing that incremental utility is inversely proportional to current fortune (and directly proportional to the increment in fortune), Bernoulli concluded that utility is a linear function of the logarithm of monetary price, and showed that in this case the moral expectation of the game is finite." (Zabell 1990, p. 13) As we shall see, Bernoulli made a mistake in the details. However, the idea that utility of wealth is not linear in wealth is the major contribution.

Various forms of the utility of wealth function are of historical interest, including the following:

the quadratic function: $U(x) = a + bx - cx^2$
for $a, b, c > 0$;

the logarithmic function: $U(x) = \log(x+b)$
for $b > 1$;

the simple exponential function: $U(x) = -ae^{-bx}$
for $a, b > 0$;

the related exponential function $U(x) = ab^{-x}$
for $a > 0, b > 1$, or $a < 0, 0 < b < 1$;

the normed exponential function $U(x) = a(1-e^{-bx})$
for $a, b > 0$; and

the normed logistic function: $U(x) = [e^{a+bx}]/[1+e^{a+bx}]$
for $a, b > 0$.

Most of these functions are flawed, and that only the latter two make sense with respect to the analysis of risk preference.

The primary attitudes toward risk are risk neutrality, risk aversion, and risk preference. The decision maker is risk neutral iff the utility function U is linear, i.e., iff $U'(x) > 0$ and $U''(x) = 0$ for all x . The decision maker is risk averse iff the utility function U is concave, i.e., iff $U'(x) > 0$ and $U''(x) < 0$ for all x . Finally, the decision maker is risk preferring iff the utility function U is convex, i.e., iff $U'(x) > 0$ and $U''(x) > 0$ for all x .

Risk aversion is generally viewed to be the most common attitude toward risk. The most recent work on utility theory has been based on experimental evidence. It focuses on testing the formal account of utility and the related decision- and game-theoretic models of which it is a part. This work has revealed much about the shape of utility functions. In particular, Fishburn and Kochenberger (1979) and Kahneman and Tversky (1979) established that individuals define utility over the *change* in wealth rather than the *amount* of wealth. Furthermore, they have shown that the reference point from which the change in wealth is calculated plays a fundamental role in decision making. The work of the psychologists and experimental economists produced the empirical evidence that most individuals have S-shaped utility functions. That is, if the decision maker's utility function is denoted by $U(x)$, where x represents change in wealth, and U' and U'' represent the first and second

derivatives, then $U(0) = 0$, $U'(x) > 0$ for all x , $U''(x) > 0$ if $x < 0$, $U''(x) < 0$ if $x > 0$, and $U'(-x) > U'(x)$ for all $x > 0$.

Specification of a Utility Function

The traditional method of specifying a utility function for a given individual is based on a reference lottery, commonly called a Becker-DeGroot-Marschak lottery. Consider a simple gamble where the decision maker receives H if a particular event occurs or receives L if the event does not occur, or receives M for sure, where $L < M < H$. Suppose the probability that the event occurs is p . We can assign a utility to M if we know the value of p . We do so as follows. Since a utility scale, like a temperature scale, allows the arbitrary selection of both the zero and the unit, we can arbitrarily set the utility values for two of the three payoffs. Common sense suggest setting the values for L and H .

Let $U(H) = 100$ and $U(L) = 0$. Then the expected utilities of the two acts are:

$$E[U(\text{take the gamble})] = pU(H) + (1-p)U(L) = p(100) + (1-p)(0) = 100p$$

and

$$E[U(\text{take the sure thing})] = U(M).$$

If we manipulate the value of p so that the decision maker is indifferent between the gamble and the sure thing, then the expected utilities are equal, and we have $U(M) = 100p$. Since we know p , we know $U(M)$.

We can continue this process for other payoffs that fall between L and H , and thereby determine enough points on the utility function to get a very good estimate of the equation of the function.

The B-D-M lottery treated here has the following payoff table:

PAYOFFS	a_1	a_2
s_1	H	M
s_2	L	M

The utility table for the lottery is then:

UTILITY	a_1	a_2
s_1	100	100p
s_2	0	100p

The Economic Theory of Information with Utility

If we introduce a nonlinear utility function, then we must adapt the approach to determining the value of information. Suppose the payoffs are as before

PAYOFFS	a_1	a_2
s_1	100	0
s_2	-50	20

And suppose the information system is as before

	P(s)	P(z/s)		P(s/z)	
		z_1	z_2	z_1	z_2
s_1	.5	.8	.2	.4/.6	.1/.4
s_2	.5	.4	.6	.2/.6	.3/.4

In order to determine the Bayes strategy and the value of the information system, we must first convert the payoff table to a utility table. Suppose the utility function is

$$U(x) = 1 - e^{-0.1x}.$$

Then the utility table is as follows:

UTILITY	a_1	a_2
s_1	1.00	0.00
s_2	-147.41	0.86

The determination of the Bayes strategy is as before, except now we use the utility values

	P(s)	P(z/s)		P(s/z)	
		z ₁	z ₂	z ₁	z ₂
s ₁	.5	.8	.2	.4/.6	.1/.4
s ₂	.5	.4	.6	.2/.6	.3/.4
				E[V(a)/z]	
				z ₁	z ₂
a ₁				-29.08/.6	-44.12/.4
a ₂				.17/.6	.26/.4

Note that the Bayes strategy is $\langle a_2, a_2 \rangle$. Thus, the decision maker's aversion to risk is sufficiently severe that even given signal z_1 , which indicates that state s_1 will occur, the decision maker selects the less risky act.

APPENDIX

The following system is an adaptation of the original von Neumann - Morgenstern system, and is taken from the textbook Mathematical Psychology by Coombs, Dawes, and Tversky, pp. 122-126.

An Axiom System for Utility Theory

Let O be the set of outcomes with typical elements $x, y, z,$ and w . A gamble is a mixture of outcomes and probabilities, and is denoted $[x, p, y]$ if outcome x is obtained with probability p and outcome y is obtained with probability $1-p$. O^* is the augmented set of outcomes and contains all of the outcomes contained in O and all of the gambles on those outcomes.

The axioms concern the system (O^*, \geq^*) where O^* is the augmented set of outcomes and \geq^* is an ordering on O^* . The ordering \geq^* supports two other orderings, $=^*$ and $>^*$, as follows:

$$\begin{aligned}x &=^* y \text{ if and only if } x \geq^* y \text{ and } y \geq^* x, \text{ and} \\x &>^* y \text{ if and only if } x \geq^* y \text{ and not } y \geq^* x.\end{aligned}$$

The axioms presented below support two theorems. The representation theorem asserts that the ordering \geq^* on O^* can be represented by a real valued function; the uniqueness theorem asserts that the real valued representation is an interval scale.

The Axioms

- Axiom 1. If x and y are elements of O , then for all probabilities $p, 0 < p < 1$, the gamble $[x, p, y]$ is an element of O^* .
- Axiom 2. \geq^* is a weak order on O^* , i.e.,
- (i) for all x in O^* , $x \geq^* x$,
 - (ii) for all x and y in O^* , $x \geq^* y$ or $y \geq^* x$ or $x =^* y$,
 - (iii) for all $x, y,$ and z in O^* , if $x \geq^* y$ and $y \geq^* z$, then $x \geq^* z$.
- Axiom 3. For all x and y in O^* and all probabilities p and $q, 0 < p, q < 1$, $[[x, p, y], q, y] =^* [x, pq, y]$.

Axiom 4. For all $x, y,$ and z in O^* and all probabilities $p, 0 < p < 1,$ if $x \sim^* y,$ then $[x, p, z] \sim^* [y, p, z].$

Axiom 5. For all x and y in O^* and all probabilities $p, 0 < p < 1,$ if $x \succ^* y,$ then $x \succ^* [x, p, y] \succ^* y.$

Axiom 6. For all $x, y,$ and z in $O^*,$ if $x \succ^* y \succ^* z,$ then there exists a probability $p, 0 < p < 1,$ such that $y \sim^* [x, p, z].$

The Representation Theorem

If the system (O^*, \succeq^*) satisfies axioms 1 through 6, then there exists a function U from (O^*, \succeq^*) to (Reals, \geq) such that

- (i) $x \succeq^* y$ if and only if $U(x) \geq U(y),$ and
- (ii) $U([x, p, y]) = pU(x) + (1-p)U(y).$

The Uniqueness Theorem

If U and V are two functions satisfying (i) and (ii) of the Representation Theorem, then there exist real numbers a and b such that $a > 0$ and for all x in $O^*, V(x) = aU(x) + b.$

Proofs of these theorems are available in von Neumann and Morgenstern (1947), Savage (1954), and elsewhere.

A Note on the Uniqueness Theorem

The uniqueness theorem for utility functions states that if an individual's preference structure can be represented by two utility functions, say $U(x)$ and $V(x),$ then U and V are related by a positive affine transformation.

(Note that $aU(x) + b$ is not a linear function; $aU(x)$ is a linear function.)

That is, there exist two numbers a and $b,$ with $a > 0,$ such that for $x, V(x) = aU(x) + b.$ You have seen this kind of a transformation before. It was used to transform temperature scales. Specifically, if $F(x)$ is the temperature of x measured on the Fahrenheit scale and $C(x)$ is the temperature measured on the centigrade scale, then $F(x) = (9/5)C(x) + 32.$ Compare this to the measurement of distance on two separate scales, say inches and centimeters. If $i(x)$ is the length of a line in inches, and $c(x)$ is the length of the same line in centimeters, then $i(x) = 2.54c(x).$

The key point is to note that the temperature conversion equation involves an additive constant, here 32, whereas the length conversion equation does not involve any additive constant. The difference rests with the arbitrary choice of the number assigned as the "zero." The "zero" on the Fahrenheit scale is 32, whereas on the Centigrade scale the "zero" is 0. On both the inch and centimeter scales, the "zero" is 0.

REFERENCES

There are numerous books on decision theory. The annotated bibliography given below includes books and other items of likely interest to managers.

Bayes, Thomas, 1763, "An Essay Towards Solving a Problem in the Doctrine of Chances," Philosophical Transactions of the Royal Society of London, Vol. 53, pp. 370-418. [Reprinted in Biometrika, Vol. 45, 1970, pp. 293-315.]

This is the paper in which Bayes presented his now famous theorem.

Bernoulli, Daniel, 1738/1954, "Exposition of a New Theory on the Measurement of Risk," Econometrica, 22, January 1954, pp. 23-36. (This is a translation from the Latin by Louise Sommer of "Specimen Theoriae Novae de Mensura Sortis," Commentarii Academiae Scientiarum Imperialis Petropolitanae, Tomus V. [Papers of the Imperial Academy of Sciences in Petersburg, Vol. 5], 1738, pp. 175-192.)

This is the original paper on accounting for risk attitude via utility function.

Bernstein, Peter L., 1996, Against the Gods: The Remarkable Story of Risk, New York: Wiley and Sons.

This is a popular history of risk.

Clemen, Robert T., 1990, Making Hard Decisions, Belmont, CA: Duxbury Press.

This is the most popular textbook on decision theory.

Fishburn, P. and Kochenberger, G., 1979, "Two-Piece von Neumann-Morgenstern Utility Functions," Decision Sciences 10, 503-518.

This paper shows that the S-shaped utility function is quite common.

Hacking, Ian, 1975, The Emergence of Probability, Cambridge: Cambridge University Press.

Hacking, Ian, 1990, The Taming of Chance, Cambridge: Cambridge University Press.

These books constitute the primary philosophical history of probability.

Hammond, John S., Ralph L. Keeney, and Howard Raiffa, 1999, Smart Choices: A Practical Guide to Making Better Decisions, Boston: Harvard Business School Press.

This is a modern and informal account of decision making.

Jeffrey, Richard C., 1983, The Logic of Decision, Second Ed., Chicago: University of Chicago Press. (The first edition was published in 1965.)

This is an elegant account of decision theory by the foremost philosopher working in the area.

Kahneman, Daniel and Amos Tversky, 1979, "Prospect Theory: An Analysis of Decisions Under Risk," Econometrica, 47, pp. 263-291.

This is the theoretical paper culminating from psychological experiments involving human decision makers..

Luce, R. Duncan and Howard Raiffa, 1957, Games and Decisions, New York: Wiley and Sons.

This is the first of the modern systematic presentations of decision theory.

Marschak, Jacob, 1974, Economic Information, Decision, and Prediction (Volume II), Boston: D. Reidel Publishing Co.

This book contains the main developmental papers by Marschak on the economic theory of information.

Ramsey, Frank P, 1931, "Truth and Probability," in The Foundations of Mathematics and Other Logical Essays, London: K. Paul, Trench, Trubner and Co.

This is one of the original papers on subjective probability.

Savage, Leonard J., 1954, The Foundations of Statistics, New York : J. Wiley and Co.

This is the primary exposition of the subjectivist account of decision making.

von Neumann, John and Oskar Morgenstern, 1947, Theory of Games and Economic Behavior, second edition, Princeton: Princeton University Press. (The first edition was published in 1944.)

This is the first modern unified account of decision making.

Zabell, S. L., 1990, "Daniel Bernoulli," pages 12-14 in J. Eatwell, M. Milgate, and P. Newman, eds., The New Palgrave: Utility and Probability, New York : W. W. Norton & Company.