Decision Theory

Refail N. Kasimbeyli

Chapter 2

2 Choosing a Decision Rule

2.1 Decision Making Under Certainty

2.2 Decision Making Under Uncertainty and Under Risk

2.3 Decision Rules forMaking Decisions Under Uncertainty

2.4 Decision Rules forMaking Decisions Under Risk

2.5 A food Manufacturer Example

Choosing a Decision Rule

An integral part of a decision problem DP of any kind is the statement of the decision rule (DR). The choice of DRs is not a trivial matter for it depends on the decision maker’s (DMr’s) attitude or the norms or policies of the governing organization. We present several common decision rules for different types of decision problems.

In the above example (see Example in Section 1.5), it is implicitly assumed that the probability distributions of the states of nature can be completely determined either subjectively or from available data.

In practice, the degree of difficulty in estimating such probabilities ranges from very trivial (e.g., in the case of certainty), to simple (e.g., tossing a coin or throwing a die), to complex (e.g., those probabilities which have to be determined subjectively), and finally, to so complex that it may not be worthwhile, or possible, to determine probabilities.

The range of difficulty and the manner in which these probabilities are estimated lead to the classification of decision problems for decision analysis, each type of which requires special solution approaches and decision rules.

2.1 Decision Making Under Certainty

Decision making under certainty corresponds to the case where it is known with certainty that one and only one of *r* states of nature occurs. In particular, that is the case in which $P(s_j) = 1$ or 0 for all $j = 1, ..., r$ and $\sum_{j=1}^{r} P(s_j) = 1$, reducing the decision problem to a deterministic one. This is obviously a special case of the standard formulation given earlier.

In a payoff matrix with only one column, if each payoff value truly represents the decision maker’s preference for the corresponding alternative, an obvious choice of decision rule would be to choose an alternative which maximizes the payoff.

In other words, this decision problem under certainty becomes a simple optimization problem of the form

$$max_{1 \leq i \leq r} u_i ,$$

where u_i is the payoff for alternative a_i. In the above example (see Example in Section 1.5), if it is known for certain that the level of demand is going to rise at the projected rate, i.e., $P(s_1) = 1$ and $P(s_2) = 0$, the decision problem would reduce to

$$max_{1 \leq i \leq r} u_{i1},$$

where solution is a_3 (purchase additional machines).

In some instances, the number of available alternatives may be so numerous (or even infinite) that it is not possible to write out the payoff matrix completely. Consequently, direct search or complete enumeration, such as above, may not be possible; more advanced techniques, such as linear programming, branch and bound, and others may have to be used. In other instances, it may not be practical to perform an extensive search for the optimal alternative.

A *satisficing* DR, in which we merely search for an alternative yielding the payoff above a prespecified threshold, would perhaps be more appropriate. Many individual decisions such as buying a house or looking for a job are of this type.

2.2 Decision Making Under Uncertainty and Under Risk

In many decisions the consequences of the alternative courses of action cannot be predicted with certainty. A company which is considering the launch of a new product will be uncertain about how successful the product will be, while an investor in the stock market will generally be unsure about the returns which will be generated if a particular investment is chosen.

In this chapter we will show how the ideas about probability, can be applied to problems where a decision has to be made under conditions of uncertainty.

We will first outline a method which assumes that the decision maker is unable, or unwilling, to estimate probabilities for the outcomes of the decision and which, in consequence, makes extremely pessimistic assumptions about these outcomes.

Then, assuming that probabilities can be assessed, we will consider an approach based on the expected value concept. Because an expected value can be regarded as an average outcome if a process is repeated a large number of times, this approach is arguably most relevant to situations where a decision is made repeatedly over a long period.

A daily decision by a retailer on how many items to have available for sale might be an example of this sort of decision problem. In many situations, however, the decision is not made repeatedly, and the decision maker may only have one opportunity to choose the best course of action.

If things go wrong then there will be no chance of recovering losses in future repetitions of the decision. In these circumstances some people might prefer the least risky course of action, and we will discuss how a decision makers attitude to risk can be assessed and incorporated into a decision model.

Decision making under uncertainty is a class of decision problems where it is not possible to estimate the probabilities of the states of nature. In such a problem, risk can no longer be quantified or analyzed explicitly, although the element of risk remains.

The most common way of coping with risk in this case is through proper selection of decision rules. Some of these rules are of the risk-prone type while others may be of either the risk-aversion or the risk-neutral type. More will be said about these rules subsequently.

Decision making under risk arises when it is possible to estimate the probability of occurrence of each state of nature $P(s_j)$ either subjectively or objectively. As the name implies, the risk can be quantified and analyzed explicitly, if we so desire. In many instances, the estimation of $P(s_j)$ for each s_j is carried out based on past data or on a purely subjective basis.

During the decision-analysis process, additional information about the environment aimed at improving the original estimates of $P(s_j)$ may be available. The utility, cost, and procedures for synthesizing additional information in this manner are the subject of Bayesian decision analysis, and the updated probabilities so constructed are called a posteriori probabilities.

2.3 Decision Rules for Making Decisions Under Uncertainty

When each $P(s_j)$ is unknown, it is not possible to use any of the maximization or minimization of expectation criteria nor is it possible to analyze risk explicitly. The following decision rules are common for this situation. They treat risk implicitly by reflecting the decision maker's attitude or intuitive feeling toward risk.

The pessimistic rule (maximin criterion) takes a pessimistic view of the environment: it assumes that no matter what alternative action a_i is selected, the worst situation for that alternative action will prevail.

A natural case of action would be to make sure that the largest possible payoff or utility under such circumstances is assured by maximizing the minimum payoff or utility (hence the name ”maximin”). Thus we solve

$$max_i \, (min_j \, u_{ij}).$$

The pessimistic view and the maximin procedure reflect risk-aversion behavior of the decision maker. Applying the maximin rule to the payoff table (Table 1.2) of the example given in Section 1.5, where $P(s_1)$ and $P(s_2)$ are now assumed to be unknown, yields the following:

$a_1 \ : \quad min_j \quad u_{1j} = 1.4 \quad , \quad \leftarrow$ maximum

$a_2 \ : \quad min_j \quad u_{2j} = 1.4 \quad , \quad \leftarrow$ maximum

$a_3 \ : \quad min_j \quad u_{3j} = 1.0 \quad .$

Hence, $max_i\,(min_j \quad u_{ij}) = 1.4$, which corresponds to either alternative a_1 or a_2 .

The optimistic rule (maximax criterion), on the other extreme, takes an optimistic view of the environment: it assumes that whatever action is taken, the maximum possible payoff or utility will result. Under such an assumption, it is natural to maximize the maximum payoff, i.e.,

$$max_i \, (max_j \, u_{ij}).$$

This is obviously a risk-prone situation. Applying the optimistic rule to the example yields

$$a_1 : \quad max_j \quad u_{1j} = 1.5 \quad ,$$

$$a_2 : \quad max_j \quad u_{2j} = 2.0 \quad ,$$

$$a_3 : \quad max_j \quad u_{3j} = 2.1 \quad . \qquad \leftarrow \text{maximum}$$

Hence, $max_i\,(max_j \quad u_{ij}) = 2.1$, which corresponds to alternative a_3.

The Hurwitz rule (optimistic index-α criterion) compromises between the two extreme viewpoints through the use of the index α. The DMaker will express his degree of optimism by specifying the value of α between 0 and 1, forming a linear (convex) combination between the maximin and maximax criteria for each alternative a_i:

$$u_i(\alpha) = \alpha min_j u_{ij} + (1 - \alpha) max_j u_{ij} \qquad (0 < \alpha < 1).$$

For a given value of α, the alternative yielding the maximum $u_i(\alpha)$ is then selected. That is we solve

$$max_i \, u_i(\alpha).$$

Applying the Hurwitz rule to our example yields

$$a_1 : \quad u_1(\alpha) = 1.4\alpha + 1.5(1 - \alpha) \qquad = 1.5 - 0.1\alpha,$$

$$a_2 : \quad u_2(\alpha) = 1.4\alpha + 2.0(1 - \alpha) \qquad = 2.0 - 0.6\alpha,$$

$$a_3 : \quad u_3(\alpha) = 1.0\alpha + 2.1(1 - \alpha) \qquad = 2.1 - 1.1\alpha.$$

Hence, we would select a_2, if $0.2 \leq \alpha \leq 1$, but would select a_3, if $0 \leq \alpha \leq 0.2$. We would never select a_1 irrespective of the value of α.

The minimax regret rule is a variation of the pessimistic rule which takes opportunity loss or regret into account rather than gain or payoff. The maximum possible loss for a given action is assumed. It would be wise, therefore to minimize the maximum opportunity loss in these circumstances. In other words, if $M_j = max_i u_{ij}$, we solve

$$min_i [max_j (M_j - u_{ij})].$$

Applying this rule to the regret table (Table 2.1) of the example again yields the selection of a_2 , since

$$a_1 : \quad max_j \quad (M_j - u_{1j}) = 0.6 \quad ,$$

$$a_2 : \quad max_j \quad (M_j - u_{2j}) = 0.1 \quad , \quad \leftarrow minimum$$

$$a_3 : \quad max_j \quad (M_j - u_{3j}) = 0.4 \quad .$$

2.4 Decision Rules for Making Decisions Under Risk

The most common decision rule for making decisions under risk is to maximize the expected payoff value. If the payoff is expressed in monetary terms, the corresponding DR would be to maximize the expected monetary value (EMV), which leads to

$$max_{1 \leq i \leq r} \sum_{j=1}^{q} P(s_j) u_{ij}.$$

In our example, the maximum-EMV decision rule leads to the selection of alternative a_2 (overtime), as we have shown earlier.

Instead of EMV, we may also use the *expected opportunity loss* (EOL), also known as *expected regret*, as our decision criterion (DC). An opportunity loss (or regret) is defined as the loss incurred due to failure to select the best alternative available. One can easily convert the payoff table to a regret table by replacing each entry of the payoff table by $M_j - u_{ij}$, where $M_j = max_{1 \leq i \leq r} u_{ij}$.

Thus in our example, the corresponding regret table is shown in Table 2.1.

Table 2.1: Regret Table for the Example

	Level of demand (millions of $)	
Action	$s_1[P(s_1) = 0.75]$	$s_2[P(s_2) = 0.25]$
a_1	0.6	0
a_2	0.1	0
a_3	0	0.4

Using the minimization of EOL decision rule leads to the selection of a_2, which is the same as when maximizing EMV.

The value judgement implied in using the EMV (or EOL) decision criterion (DC) would be that the higher the monetary payoff (or the lower the value of opportunity loss) we can expect, the better the alternative action. Note that utility is always assumed to be a monotonic function and that the attitude toward risk is invariant with respect to the size of the payoff.

Such a premise is, however, frequently subjected to criticism since the element of risk is not considered. One effective way to incorporate risk in the analysis is through explicit quantification of risk.

To be specific, we define risk associated with choosing an alternative a_i as the probability that the actual payoff falls below a prespecified level β given that an alternative a_i is selected. Denote this probability by ρ_i. Clearly, ρ_i can be easily computed from the formula

$$\rho_i = \sum_{j \in J_i} P(s_j) \quad \text{where} \quad J_i = \{j : 1 \leq j \leq n, u_{ij} \leq \beta\}$$

By convention, $\rho_i = 0$ if J_i is empty (i.e., all $u_{ij} \geq \beta$), in which case there is no risk that the payoff will fall below a given value if alternative a_i is chosen. The probability ρ_i can be computed for each alternative.

We can then modify the EMV DR to read, ”Choose any alternative that maximizes EMV and has ρ_i no greater than a prespecified risk level α,” which solves the following problem:

$$\begin{array}{ll} max_{1\leq i\leq r} & \sum_{j=1}^{q} P(s_j)u_{ij} \\ \text{subject to} & \rho_i \leq \alpha \end{array}.$$

In our example, suppose we set the lower limit of payoff β to be \$1.2 million and the highest permissible level of risk at $\alpha = 0.20$. Clearly, alternative a_3 would not qualify, despite the fact that it has the highest possible payoff value, since $\rho_3 = 0.25 > \alpha$. Both a_1 and a_2 are possible candidates since $\rho_1 = \rho_2 = 0$. But EMV of a_2 is greater than that of a_1, which leads to the selection of a_2 according to the modified DR.

A different way of treating risk is by means of another form of value judgement. In this approach, monetary values are replaced by ***utility***, which measures the decision maker's preference in an interval scale.

Using utility as a measure of reward or gain, the corresponding decision rule would be to maximize the expected value of utility (EUV). The rationale of this approach is that the same amount of money may be valued differently by different individuals. under different circumstances.

A classic example that will help illustrate this point is a simple coin-tossing game.

Suppose B challenges A to play the following simple game. A fair coin is tossed once. If a head turns up, B will pay A \$400, and if tails turn up, A must pay B \$200. Should A play the game? If A is an EMVer, he should clearly play the game since EMV for playing is

$$\frac{1}{2}(\$400) + \frac{1}{2}(-\$200) = \$100,$$

which is greater than EMV for not playing (namely zero).

However, there is also a 50% chance that *A* will loss $200. If *A* is a risk-conscious individual, then he would consider the value of $200 against his financial status. If the loss of $200 has minimal effect on his finances, then he would probably play the game. Or in the other hand, *A* may have only a small bank account, which he put aside for emergencies. He may want to avoid any risk of losing $200, even with an EMV of $100.

The theoretical basis for representing preference by utility will be discussed in the next section.

2.5 A food Manufacturer Example

Consider the following problem. Each morning a food manufacturer has to make a decision on the number of batches of a perishable product which should be produced. Each batch produced costs \$800 while each batch sold earns revenue of \$1000.

Any batch which is unsold at the end of the day is worthless. The daily demand for the product is either one or two batches, but at the time of production the demand for the day is unknown and the food manufacturer feels unable to estimate probabilities for the two levels of demand.

The manufacturer would like to determine the optimum number of batches which he should produce each morning. Clearly the manufacturer has a dilemma. If he produces too many batches, he will have wasted money in producing food which has to be destroyed at the end of the day. If he produces too few, he will be forgoing potential profits. We can represent his problem in the form of a Table 2.2.

Table 2.2: Data related to the food manufacturer example

	(Daily Profits)	Demand (no. of batches)
Course of action	1	2
Produce 1 batch	\$200	\$200
Produce 2 batches	–\$600	\$400

The rows of this table represent the alternative courses of action which are open to the decision maker (i.e. produce one or two batches), while the columns represent the possible levels of demand which are, of course, outside the control of the decision maker. The monetary values in the table show the profits which would be earned per day for the different levels of production and demand.

For example, if one batch is produced and one batch demanded, a profit of \$1000 - \$800 (i.e. \$200) will be made. This profit would also apply if two batches were demanded, since a profit can only be made on the batch produced.

Given these potential profits and losses, how should the manufacturer make his decision? (We will assume that he has only one objective, namely maximizing monetary gain so that other possible objectives, such as maximizing customer goodwill or market share, are of no concern.)

Decision making under certainty: the maximin criterion

According to the maximin criterion the manufacturer should first identify the worst possible outcome for each course of action and then choose the alternative yielding the best of these worst outcomes.

If the manufacturer produces one batch, he will make the same profit whatever the demand, so the worst possible outcome is a profit of $200. If he decides to produce two batches the worst possible outcome is a loss of $600.

As shown below, the best of these worst possible outcomes (the maximum of the minimum possible profits) is associated with the production of one batch per day so, according to maximin, this should be the manufacturers decision.

Table 2.3: A decision table for the food manufacturer

Course of action	*Worst possible profit*	
Produce 1 batch	\$200	– best of the worst possible outcomes
Produce 2 batches	–\$600	

Note that if the outcomes had been expressed in terms of costs, rather than profits, we would have listed the highest possible costs of each option and selected the option for which the highest possible costs were lowest. Because we would have been selecting the option with the minimum of the maximum possible costs our decision criterion would have been referred to as *minimax*.

The main problem with the maximin criterion is its inherent pessimism. Each option is assessed only on its worst possible outcome so that all other possible outcomes are ignored. The implicit assumption is that the worst is bound to happen while, in reality, the chances of this outcome occurring may be extremely small.

For example, suppose that you were offered the choice of receiving $1 for certain or taking a gamble which had a 0.9999 probability of yielding $1 million and only a 0.0001 probability of losing you $1.

The maximin criterion would suggest that you should not take the risk of engaging in the gamble because it would assume, despite the probabilities, that you would lose.

This is unlikely to be a sensible representation of most decision makers preferences. Nevertheless, the extreme risk aversion which is implied by the maximin criterion may be appropriate where decisions involve public safety or possible irreversible damage to the environment.

A new cheaper form of food processing which had a one in ten thousand chance of killing the entire population would clearly be unacceptable to most people.

Decision making under risk: The expected monetary value (EMV) criterion

If the food manufacturer is able, and willing, to estimate probabilities for the two possible levels of demand, then it may be appropriate for him to choose the alternative which will lead to the highest expected daily profit. If he makes the decision on this basis then he is said to be using the expected monetary value or EMV criterion.

Recall that an expected value can be regarded as an average result which is obtained if a process is repeated a large number of times. This may make the criterion particularly appropriate for the retailer who will be repeating his decision day after day.

Table 2.4 shows the manufacturers decision table again, but this time with the probabilities added.

Table 2.4: Another decision table for the food manufacturer

		(Daily Profits)	Demand (no. of batches)
		1	2
Course of action	Probability	0.3	0.7
Produce 1 batch		\$200	\$200
Produce 2 batches		−\$600	\$400

An expected value is calculated by multiplying each outcome by its probability of occurrence and then summing the resulting products.

The expected daily profits for the two production levels are therefore:

Produce one batch:

expected daily profit $=(0.3 \times \$200)+(0.7 \times \$200)=\$200$

Produce two batches:

expected daily profit $=(0.3\times(-\$600))+(0.7\times\$400)=\$100$.

These expected profits show that, in the long run, the highest average daily profit will be achieved by producing just one batch per day and, if the EMV criterion is acceptable to the food manufacturer, then this is what he should do.

Of course, the probabilities and profits used in this problem may only be rough estimates or, if they are based on reliable past data, they may be subject to change. We should therefore carry out sensitivity analysis to determine how large a change there would need to be in these values before the alternative course of action would be preferred.

To illustrate the process, Figure 2.1 shows the results of a sensitivity analysis on the probability that just one batch will be demanded. Producing one batch will always yield an expected profit of $200, whatever this probability is.

However, if the probability of just one batch being demanded is zero, then the expected profit of producing two batches will be \$400. At the other extreme, if the probability of just one batch being demanded is 1.0 then producing two batches will yield an expected profit of −\$600.

The line joining these points shows the expected profits for all the intermediate probabilities. It can be seen that producing one batch will continue to yield the highest expected profit as long as the probability of just one batch being demanded is greater than 0.2.

Since currently this probability is estimated to be 0.3, it would take only a small change in the estimate for the alternative course of action to be preferred. Therefore in this case the probability needs to be estimated with care.

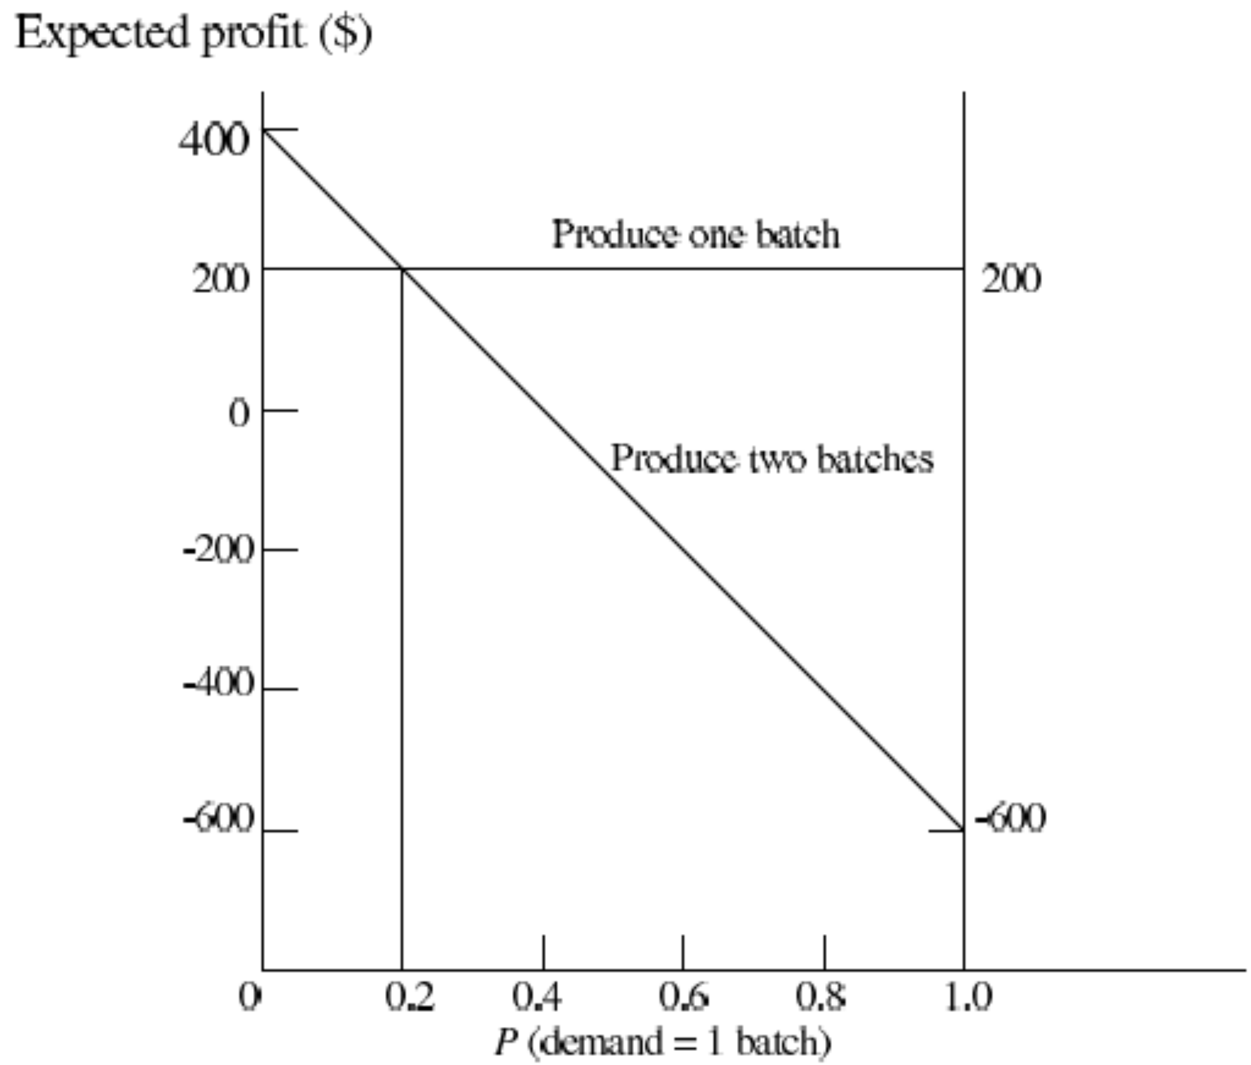

Figure 2.1: A sensitivity analysis for the food manufacturer's problem.