

3 Decision Theory

3.1 Utility Functions

Suppose the utility of having x units of a good is $u(x)$.

It is normally assumed that u is a concave, increasing function and $u(0) = 0$.

i.e. the utility from having $2x$ Euros is greater than the utility of having x Euros, but not more than twice as great.

It is assumed that individuals act so as to maximise their (expected) utility.

Utility functions

Utility functions for the acquisition of n goods can be defined in a similar way.

Suppose $\underline{x} = (x_1, x_2, \dots, x_n)$ defines a bundle of goods, i.e. an individual has x_i units of good i .

The utility from having such a bundle is defined to be $u(\underline{x})$, where $u(\underline{0}) = 0$, $\frac{\partial u}{\partial x_i} > 0$ and $\frac{\partial^2 u}{\partial x_i \partial x_j} \leq 0$.

Utility functions and optimality in deterministic problems

From the assumptions made, if a deterministic problem only involves the acquisition of one good (commonly money in the form of profit), maximising profit will automatically maximise utility.

Hence, the form of the utility function will not have any effect on the optimal solution of such a problem.

Utility functions and optimality in deterministic problems

However, if a problem involves the acquisition of two goods, then the form of the utility function will in general have an effect on the optimal solution,

e.g. when allowed to choose five pieces of fruit some individuals may prefer to take 3 apples and 2 oranges, while others prefer 2 oranges and 3 apples.

However taking 2 apples and 2 oranges would clearly be a sub-optimal solution.

Concavity of utility functions

The assumption that a utility function is concave seems reasonable as follows:

The utility I gain by obtaining an apple (the marginal utility of an apple) when I have no apples is greater than the utility I gain by obtaining an apple when I already have a large number of apples.

It should be noted that

- a) Individuals have different utility functions.
- b) A person's utility function depends on his/her present circumstances.

Utility functions and optimality in probabilistic problems

In probabilistic models we maximise expected utility. In general, the optimal solution depends on the form of this utility function.

The concept of utility is related to **Jensen's inequalities**.

Jensen's Inequalities

If f is a concave function, g is a convex function and X a random variable then

$$E[f(X)] \leq f[E(X)]$$

$$E[g(X)] \geq g[E(X)]$$

Risk neutral individuals

Suppose $u(x)$ is an individual's utility function for obtaining x units of money.

If $u''(x) = 0$ for all x (i.e. $u(x)$ is linear), then that individual is said to be risk neutral.

When presented with the choice between a fixed sum of money k and lottery in which the expected prize is k , such an individual is indifferent between these choices.

Risk averse individuals

If a person has a utility function that is strictly concave, then that individual is said to be risk averse.

When presented with the choice between a fixed sum of money k and lottery in which the expected prize is k , such an individual would choose the fixed sum.

This follows from the fact that if u is a strictly concave function $E[u(X)] < u[E(X)]$. The first expression is the expected utility from playing the lottery. The second expression is the utility from choosing the fixed sum.

Example 3.1

Consider the following situation. A contestant in a game show must choose between

- a) A guaranteed reward of 125 000 Euros.
- b) Taking one of two bags. One of the bags is empty, the other contains 250 000 Euro.

It can be seen that by taking option b), the expected amount of money obtained by the player is also 125 000 Euros.

Example 3.1

We consider 3 utility functions describing the utility of obtaining x Euro

$$u_1(x) = x; \quad u_2(x) = \sqrt{x}; \quad u_3(x) = x^2$$

Example 3.1 - Risk neutral individual

If an individual has a utility function given by $u_1(x) = x$, then by taking option a), he/she obtains an expected utility of 125 000.

By taking option b), he/she obtains an expected utility of

$$E[u_1(X)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 250000 = 125000.$$

Hence, this player is indifferent between choosing a guaranteed amount of money and a lottery which gives the same expected amount of money (i.e. is risk neutral).

Example 3.1 - Risk averse individual

If an individual has a utility function given by $u_2(x) = \sqrt{x}$, then by taking option a), he/she obtains an expected utility of $\sqrt{125000} \approx 353.55$.

By taking option b), he/she obtains an expected utility of

$$E[u_2(X)] = \frac{1}{2} \times 0 + \frac{1}{2} \times \sqrt{250000} = 250.$$

Hence, this player would rather choose a guaranteed amount of money than a lottery which gives the same expected amount of money (i.e. is risk averse).

Example 3.1 - Risk averse individual

We can calculate the guaranteed amount x_1 for which the player is indifferent between that amount and the lottery as follows:

$$u_2(x_1) = \sqrt{x_1} = 250 \Rightarrow x_1 = 62500.$$

Example 3.1 - Risk seeking individual

If an individual has a utility function given by $u_3(x) = x^2$, then by taking option a), he/she obtains an expected utility of $125000^2 = 1.625 \times 10^{10}$.

By taking option b), he/she obtains an expected utility of

$$E[u_3(X)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 250000^2 = 3.125 \times 10^{10}.$$

Hence, this player would rather choose a lottery than a guaranteed amount of money, if the lottery gives the same expected amount of money (i.e. is risk seeking).

Decision Analysis - Deterministic models

When making a decision we must often take various criteria into account.

For example, when I want to buy an airline ticket I do not just take the price of the ticket into account.

Other factors I take into account may be the time of the flight, the distance of airports to a) my home b) to my destination etc.

Decision Analysis - Deterministic models

In order to decide on which ticket to buy, I may ascribe

- i) Various weights to these factors.
- ii) Scores describing the attractiveness of price, time of flights and location.

Such an approach is commonly used in defining the bid a public sector company should accept.

Decision Analysis - Deterministic models

I choose the option that maximises my utility, which is defined to be the appropriate weighted average of these scores.

In mathematical terms, suppose there are k factors. Let w_i be the weight associated with factor i and $s_{i,j}$ the score of option j according to factor i (i.e. how good option j is according to factor i).

It is normally assumed that $\sum_{i=1}^k w_i = 1$.

Decision Analysis - Deterministic models

I choose the option j , which maximises the utility of option j , u_j , which is taken to be the weighted average of the scores i.e.

$$u_j = \sum_{i=1}^k w_i s_{i,j}$$

Hence, a mathematical model of such a problem is defined by

- a) The set of options.
- b) The set of factors influencing the decision.
- c) Weights for each of these factors.
- d) Scores describing how attractive each option is according to each factor.

Example 3.2

Suppose I am travelling to Barcelona and I consider two possibilities.

I can fly using a cheap airline from Shannon to Girona (about 70km from Barcelona), with one of the flights being at a very early time or

I can fly using a more expensive airline from Shannon to Barcelona and the times of the flight are convenient.

Example 3.2

The three factors I consider are price (factor 1), location (factor 2) and time (factor 3). I ascribe the weights $w_1 = 0.5$, $w_2 = 0.3$ and $w_3 = 0.2$ to each of these factors.

The budget airline scores $s_{1,1} = 90$ according to price, $s_{2,1} = 40$ according to location and $s_{3,1} = 60$ according to time.

The other airline scores $s_{1,2} = 70$ according to price, $s_{2,2} = 90$ according to location and $s_{3,2} = 90$ according to time.

Example 3.2 - Utility of the options

Hence,

$$u_1 = 0.5 \times 90 + 0.3 \times 40 + 0.2 \times 60 = 69$$

$$u_2 = 0.5 \times 70 + 0.3 \times 90 + 0.2 \times 90 = 80.$$

It follows that I should choose the flight from Shannon to Barcelona.

Multi-person decision processes

When more than one person is involved in the decision process, we may define a composite utility as a weighted average of the utilities of the alternatives to each of the decision makers.

Suppose Adam and Betty want to meet in one of two restaurants. They take two factors into account a) the location of a restaurant and b) the attractiveness of a restaurant.

Multi-person decision processes

In this case we need to define

- a) Weights describing the importance of each person in the decision process (say p and q here).
- b) The weights Adam gives to the importance of location, p_1 and attractiveness, p_2 and the scores he gives to restaurant j according to these criteria $r_{1,j}$ and $r_{2,j}$.
- c) The weights Betty gives to the importance of location, q_1 and attractiveness, q_2 and the scores she gives to restaurant j according to these criteria $s_{1,j}$ and $s_{2,j}$.

These scores must be made on the same scale.

Multi-person decision processes

Under these assumptions we can define the utility of the choice of restaurant j to both Adam, $u_{A,j}$, and Betty, $u_{B,j}$.

These are simply the weighted averages of the scores each one gives for location and attractiveness. We have

$$u_{A,j} = \sum_{i=1}^2 p_i r_{i,j}$$
$$u_{B,j} = \sum_{i=1}^2 q_i s_{i,j}$$

Multi-person decision processes

The composite utility is the weighted average of these utilities (weighted according to the importance of the decision makers).

In this case, the composite utility gained from choosing the j -th restaurant is u_j , where

$$u_j = pu_{A,j} + qu_{B,j}.$$

The decision makers should choose the restaurant which maximises the composite utility.

Multi-person decision processes

For example, suppose the decision makers are of equal importance, i.e. $p = q = 0.5$.

Adam's weights for the importance of location and attractiveness are 0.7 and 0.3.

Betty's weights are 0.4 and 0.6, respectively.

Multi-person decision processes

They must choose between 2 restaurants. Adam assigns a score of 70 (out of 100) for location and 50 for attractiveness to Restaurant 1 and a score of 40 for location and 90 for attractiveness to Restaurant 2.

The analogous scores given by Betty are 50 and 60 (to Restaurant 1) and 30 and 80 (to Restaurant 2).

Multi-person decision processes

The utilities to Adam of these restaurants are

$$u_{A,1} = 0.7 \times 70 + 0.3 \times 50 = 64$$

$$u_{A,2} = 0.7 \times 40 + 0.3 \times 90 = 55.$$

The utilities to Betty of these restaurants are given by

$$u_{B,1} = 0.4 \times 50 + 0.6 \times 60 = 56$$

$$u_{B,2} = 0.4 \times 30 + 0.6 \times 80 = 60.$$

Multi-person decision processes

Hence, the composite utilities of these choices are given by

$$u_1 = 0.5 \times 64 + 0.6 \times 56 = 60$$

$$u_2 = 0.5 \times 55 + 0.6 \times 60 = 57.5.$$

It follows that Restaurant 1 should be chosen.

Probabilistic models

In many cases we do not know what the outcome of our decisions will be, but we can ascribe some likelihood (probability) to these outcomes.

In this case we can define a decision tree to illustrate the actions we may take and the possible outcomes resulting from these actions.

Probabilistic models

Consider the following simplified model of betting on a horse.

Suppose there are only 2 horses we are interested in. If the going is good or firmer (this occurs with probability 0.6), horse 1 wins with probability 0.5, otherwise it wins with probability 0.2.

If the going is good horse 2 wins with probability 0.1, otherwise it wins with probability 0.4.

Probabilistic models

The odds on horse 1 are evens. The odds on horse 2 are 2 to 1.

Hence (ignoring taxes), if we place a 10 Euro bet on horse 1, we win 10 Euro if it wins. If we place a 10 Euro bet on horse 2, we win 20 Euro if it wins.

In all other cases we lose 10 Euro. Winnings are denoted by W .

Probabilistic models

It should be noted that if the odds are valid only before the ground conditions are known, we are only really interested whether these horses win or lose and not in whether the ground is good or not.

Hence, to simplify the decision tree, we calculate the probability that a horse wins.

Let A be the event that the first horse wins and let B be the event that the second horse wins.

Let G be the event that the ground is good or firmer (G^c is the complement of event G , i.e. not G).

Probabilistic models

Using the law of complete probability

$$P(A) = P(A|G)P(G) + P(A|G^c)P(G^c) = 0.5 \times 0.6 + 0.2 \times 0.4 = 0.38$$

$$P(B) = P(B|G)P(G) + P(B|G^c)P(G^c) = 0.1 \times 0.6 + 0.4 \times 0.4 = 0.22$$

Probabilistic models

The simplified decision tree for this problem is given by

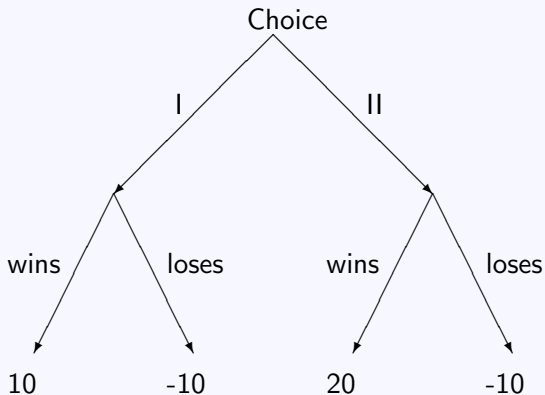


Fig. 1: Decision Tree for Betting Problem

Criteria for Choosing under Uncertainty

In the case where risk exists, there are various criteria for choosing the optimal action. The first criterion we consider is **the criterion of maximising the expected amount of the good to be obtained**. Here we should maximise the expected winnings.

This criterion is valid for individuals who are indifferent to risk (i.e. have a linear utility function).

By betting on horse 1, my expected winnings (in Euros) are

$$E(W) = 10 \times 0.38 - 10 \times 0.62 = -2.4.$$

By betting on horse 2, my expected winnings are

$$E(W) = 20 \times 0.22 - 10 \times 0.78 = -3.4.$$

Thus, it is better to bet on horse 1 than on horse 2. Of course, in this case it would be better not to bet at all.

Information in probabilistic problems

One interesting point to note regarding this example is that there is a signal (the state of the ground), which is correlated with the results of the race.

It was assumed that the bet took place before the signal could be observed.

The probabilities of each horse winning in this case are termed the "a priori" probabilities.

Information in probabilistic problems

Suppose now the signal was observed before the bet was made.

In this case we may base our bet on the signal, using the "posterior" probabilities of each horse winning.

These are the conditional probabilities of the horses winning given the state of the ground.

Hence, we should define the decision to be made when the conditions are good or firmer D_G and the decision to be made when the conditions are different D_{NG} , i.e. we define the decision to be made for each possible signal.

Information in probabilistic problems

Suppose the conditions are good or firmer. If I bet on horse 1 (who wins with probability 0.5), my expected winnings are

$$E(W) = 10 \times 0.5 - 10 \times 0.5 = 0$$

If I bet on horse 2 (who wins with probability 0.1), my expected winnings are

$$E(W) = 20 \times 0.1 - 10 \times 0.9 = -7$$

Hence, when the conditions are good or firmer I should bet on horse 1. In this case I am indifferent between betting and not betting.

Information in probabilistic problems

Suppose the conditions are softer than good. If I bet on horse 1 (who wins with probability 0.2), my expected winnings are

$$E(W) = 10 \times 0.2 - 10 \times 0.8 = -6$$

If I bet on horse 2 (who wins with probability 0.4), my expected winnings are

$$E(W) = 20 \times 0.4 - 10 \times 0.6 = 2$$

Hence, when the conditions are softer I should bet on horse 2. In this case, I prefer betting to not betting.

Maximisation of expected utility

The second criterion we consider is the **maximisation of utility**.

Although the criterion of maximisation of the expected amount of the good to be obtained is very simple to use, some people may be more averse to risk than others.

In this case it may be useful to use the maximisation of utility.

Example 3.3

Consider the previous example in which the ground conditions are observed before the bet is made.

Assume that my utility function is $u(x) = \sqrt{10 + x}$, where x are my winnings.

Note that $10 + x$ is the amount of money I have after the bet.

Example 3.3

Suppose the conditions are good or firmer. If I bet on horse 1 (who wins with probability 0.5), my expected utility is

$$E(W) = \sqrt{20} \times 0.5 + 0 \times 0.5 = 2.236$$

If I bet on horse 2 (who wins with probability 0.1), my expected winnings are

$$E(W) = \sqrt{30} \times 0.1 - 0 \times 0.9 = 0.5477$$

If I do not bet, then my utility is $\sqrt{10} = 3.162$.

Hence, when the conditions are good or firmer (if I bet) I should bet on horse 1. In this case I would rather not bet.

Example 3.3

Suppose the conditions are softer than good. If I bet on horse 1 (who wins with probability 0.2), my expected utility

$$E(W) = \sqrt{20} \times 0.2 - 0 \times 0.8 = 0.8944$$

If I bet on horse 2 (who wins with probability 0.4), my expected utility is

$$E(W) = \sqrt{30} \times 0.4 - 0 \times 0.6 = 2.191$$

Hence, when the conditions are softer (if I bet) I should bet on horse 2. However in this case, I prefer not betting to betting.

Other criteria for choice under uncertainty

Other criteria for choosing under uncertainty are

- 1) Laplace's criterion
- 2) The minimax criterion
- 3) The Savage criterion
- 4) The Hurwicz criterion

Laplace's criterion

The **Laplace criterion** is based on the **principle of insufficient reason**.

Since normally we do not know the probabilities of the events that interest us, we should assume that these are equally likely.

This means that if there are n horses in a race, then we assume that each has a probability $\frac{1}{n}$ of winning. In this case it is clear that in order to maximise our expected winnings (or utility) any bet should be placed on the horse with the largest odds.

The minimax criterion

The **minimax** criterion is based on the conservative attitude of minimising the maximum possible loss (this is equivalent to maximising the minimum possible gain).

The Savage criterion

The **Savage (regret) criterion** aims at moderating conservatism by introducing a regret matrix.

Let $v(a_i, s_j)$ be the payoff (or loss) when a decision maker takes action a_i and the state (result of the random experiment) is s_j .

If v denotes a payoff, then the measure of regret $r(a_i, s_j)$ is given by

$$r(a_i, s_j) = \max_{a_k} \{v(a_k, s_j)\} - v(a_i, s_j).$$

This is the maximum gain a decision maker could have made, if he/she knew what the state of nature was going to be.

The Savage criterion

If v denotes a loss, then the measure of regret $r(a_i, s_j)$ is given by

$$r(a_i, s_j) = v(a_i, s_j) - \min_{a_k} \{v(a_k, s_j)\}.$$

This is the maximum decrease in costs a decision maker could have made (relative to the loss he/she actually incurred), if he/she knew what the state of nature would be.

The decision is made by applying the minimax criterion to the regret matrix. Since regret is a "cost", we minimise the maximum regret.

The Hurwitz criterion

The **Hurwitz** criterion is designed to model a range of decision-making attitudes from the most conservative to the most optimistic.

Define $0 \leq \alpha \leq 1$ and suppose the gain function is $v(a_i, s_j)$.

The action selected is the action which maximises

$$\alpha \max_{s_j} v(a_i, s_j) + (1 - \alpha) \min_{s_j} v(a_i, s_j).$$

The Hurwitz criterion

The parameter α is called the **index of optimism**.

If $\alpha = 0$, this criterion is simply the minimax criterion (i.e. the minimum gain is maximised, a conservative criterion).

If $\alpha = 1$, then the criterion seeks the maximum possible payoff (an optimistic criterion). A typical value for α is 0.5.

The Hurwitz criterion

If $v(a_i, s_j)$ is a loss, then the criterion selects the action which minimises

$$\alpha \min_{s_j} v(a_i, s_j) + (1 - \alpha) \max_{s_j} v(a_i, s_j).$$

Example 3.4

Suppose a farmer can plant corn (a_1), wheat (a_2), soybeans (a_3) or use the land for grazing (a_4).

The possible states of nature are: heavy rainfall (s_1), moderate rainfall (s_2), light rainfall (s_3) or drought (s_4). The payoff matrix in thousands of Euro is as follows

	s_1	s_2	s_3	s_4
a_1	-20	60	30	-5
a_2	40	50	35	0
a_3	-50	100	45	-10
a_4	12	15	15	10

Example 3.4

Using the four criteria given above, determine the action that the farmer should take.

Use an optimality index $\alpha = 0.4$.

Example 3.4 - Using Laplace's criterion

Using Laplace's criterion, we assume each of the states is equally likely and we calculate the expected reward obtained for each action $E(W; a_i)$.

We have

$$E(W; a_1) = \frac{-20 + 60 + 30 - 5}{4} = 16.25$$

$$E(W; a_2) = \frac{40 + 50 + 35 + 0}{4} = 31.25$$

$$E(W; a_3) = \frac{-50 + 100 + 45 - 10}{4} = 21.25$$

$$E(W; a_4) = \frac{12 + 15 + 15 + 10}{4} = 13$$

Under this criterion, the farmer should take action a_2 (plant wheat).

Example 3.4 - Using the minimax criterion

Under the minimax criterion, we maximise the minimum possible reward.

Let $W_{min}(a_i)$ be the minimum possible reward given that action a_i is taken

$$W_{min}(a_1) = -20; \quad W_{min}(a_2) = 0; \quad W_{min}(a_3) = -50; \quad W_{min}(a_4) = 10.$$

Under this criterion, the farmer should take action a_4 (use for grazing).

Example 3.4 - Using the Savage criterion

When using Savage's criterion, we first need to calculate the regret matrix.

The values in the column of the regret matrix corresponding to state s_j are the gains that a decision maker could make by switching his action from the action taken to the optimal action in that state.

Given the state is s_1 , the optimal action is a_2 , which gives a payoff of 40.

Hence, the regret from playing a_1 is 60. The regret from playing a_3 is 90 and the regret from playing a_4 is 28.

Example 3.4 - Using the Savage criterion

Calculating in a similar way, the regret matrix is given by

	s_1	s_2	s_3	s_4
a_1	60	40	15	15
a_2	0	50	10	10
a_3	90	0	0	20
a_4	28	85	30	0

Example 3.4 - Using the Savage criterion

In order to choose the appropriate action, we apply the minimax criterion to the regret matrix.

Since regret is a cost (loss), we minimise the maximum regret.

Let $R_{max}(a_i)$ be the maximum regret possible given that action a_i is taken. We have

$$R_{max}(a_1) = 60; \quad R_{max}(a_2) = 50; \quad R_{max}(a_3) = 90; \quad R_{max}(a_4) = 85$$

It follows that according to this criterion, the action a_2 should be taken (plant wheat).

Example 3.4 - Using the Hurwicz criterion

Finally, we consider the Hurwicz criterion with $\alpha = 0.4$. Let the Hurwicz score of action a_1 be given by

$$H(a_i) = \alpha \max_{s_j} v(a_i, s_j) + (1 - \alpha) \min_{s_j} v(a_i, s_j).$$

If the farmer takes action a_1 , the maximum possible reward is 60 and the minimum possible reward is -20. The Hurwicz score is thus given by

$$H(a_1) = 0.4 \times 60 + 0.6 \times (-20) = 12.$$

Example 3.4 - Using the Hurwicz criterion

Similarly,

$$H(a_2) = 0.4 \times 50 + 0.6 \times 0 = 20$$

$$H(a_3) = 0.4 \times 100 + 0.6 \times (-50) = 10$$

$$H(a_4) = 0.4 \times 15 + 0.6 \times 10 = 12.$$

Hence, according to this criterion action a_2 (plant wheat) should be taken.

It can be seen that for small α (the decision maker is pessimistic), then action a_4 should be taken. If α is close to 1 (the decision maker is optimistic), then action a_3 should be taken.