Decision Theory

Simon Shaw
s.c.shaw@maths.bath.ac.uk

2005/06 Semester II

1 Introduction

We face a decision problem whenever there is a choice between two or more courses of action

> e.g. today, you all had a choice about whether to attend this lecture.

It is often not enough to include just one action and its negation. You have to consider all the possible alternatives.

> e.g. if you decide not to attend the lecture, you have to do something with your time (go shopping; go to the pub; go to the library)

Your decision problem is a choice amongst these.

1.1 Actions

The first step of the decision process is to construct a list of the possible actions that are available.

> e.g. A_1 (go to the lecture); A_2 (go shopping); A_3 (go to the pub); A_4 (go to the library)

Our theory will limit us to a selection amongst this list of actions so the list must exhaust the possibilities, i.e. it is an **exhaustive** list. We require that we will decide to take only one of the actions, that is the list is **exclusive**. Alternatively, one of the actions has to be taken, and at most one can be taken.

> e.g. you will either go to the lecture, go shopping, go to the pub or go to the library. You cannot do any two or more of these at the same time.

Definition 1 *(Action Space)*
Suppose that $A_1, A_2, \dots, A_m$ is a list of m exclusive and exhaustive actions. The set $\mathcal{A} = \{A_1, A_2, \dots, A_m\}$ is called the action space.

1.2 State of Nature

The selection of a single action as being, in some sense, best and the one course of action to adopt is straightforward providing one has complete information.

> e.g. you would choose to go to the lecture if you knew that in the lecture you were to find out exactly what was in the exam.

	S_1	S_2	S_3
A_1	88	53	20
A_2	75	66	32
A_3	57	50	39

Table 1: Payoff (profit) table for Duff beer example

The difficulty in selection is due to you not knowing exactly what will happen if a particular course of action is chosen. We are concerned with **decision making under uncertainty.** The second stage of the decision process is to identify all relevant possible uncertainties.

Definition 2 *(States of nature)*
Suppose that $S_1, S_2, \ldots, S_n$ is a list of n exclusive and exhaustive events. The set $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ is known as the states of nature or parameter space.

The decision process is to choose a single action, not knowing the state of nature.

1.3 Payoff

In order to compare each combination of action and state of nature we need a payoff (e.g. profit or loss). Typically, this will be a numerical value and it will be clear how we compare the payoffs. For example, we seek to maximise profit or to minimise loss. We shall deal with maximisation problems (note that, as one is the negation of the other, there is a duality between maximisation and minimisation problems). Initially, we shall consider the payoff to be monetary.

Definition 3 *(Payoff)*
Associated with each action and each state of nature is a payoff

$$\pi_{ij} = \{\pi(A_i, S_j), A_i \in \mathcal{A}, S_j \in \mathcal{S}\}.$$

The payoff represents the 'reward' obtained by opting for action A_i when the state of nature turns out to be S_j.

1.4 Working Example

Duff beer would like to market a new beer, 'Tartar Control Duff'. The manager is trying to decide whether to produce the beer in large quantities (A_1), moderate quantities (A_2) or small quantities (A_3). The manager does not know the demand for his product, but asserts that three events could occur: strong demand (S_1), moderate demand (S_2) or weak demand (S_3). The profit, in thousands of dollars, with regard to marketing the beer is given in the corresponding **payoff table**, see Table 1. The payoff table represents the payoffs in tabular form: the rows are the actions and the columns are the states of nature. For example, $\pi_{12} = 53$.

What should Duff beer do? If they know the state of nature then, as more profit is preferred to less, then the decision would be easy:

- if S_1 occurs, choose A_1

- if S_2 occurs, choose A_2

- if S_3 occurs, choose A_3

The choice of action which maximises profit depends upon the state of nature.

1.5 Admissibility

Sometimes, before we choose which action to take, it is possible to exclude certain actions from further consideration because they can never be the best.

Definition 4 *(Dominate; inadmissible)*
If the payoff for A_i is at least as good as that for A_j, whatever the state of nature, and is better than that for A_j for at least one state of nature then action A_i is said to dominate action A_j. Any action that is dominated is said to be inadmissible.

Inadmissible actions should be removed from further consideration. In the payoff table for Duff beer, all the actions are admissible.

> NB. if π_{21} had been greater than, or equal to, 88 then A_2 would have dominated A_1 and A_1 would have been inadmissible.

1.6 Minimax regret

Let $\pi^*(S_j)$ denote the best decision under state of nature S_j. Thus, for maximisation problems, we have

$$\pi^*(S_j) \quad = \quad \max_i \pi(A_i, S_j).$$

> i.e. the largest value in the jth column.

Definition 5 *(Opportunity regret)*
Suppose we select action A_i and learn state S_j has occurred. The difference $R(A_i, S_j)$ between the best payoff $\pi^*(S_j)$ and the actual payoff $\pi(A_i, S_j)$ is termed the opportunity loss or regret associated with action A_i when state S_j occurs.

Thus, the opportunity regret is the payoff lost due to the uncertainty in the state of nature. The regret for action A_i and state S_j is calculated as

$$R(A_i, S_j) \quad = \quad \pi^*(S_j) - \pi(A_i, S_j).$$

You wish to avoid opportunity losses. The **minimax regret strategy** is to choose the action that minimises your maximum opportunity loss.

1.6.1 Duff beer example

For the Duff beer example, $\pi^*(S_1) = 88$, $\pi^*(S_2) = 66$ and $\pi^*(S_3) = 39$. Table 2 shows the corresponding opportunity regrets. From Table 2 we find the largest opportunity regret for each action

Action	Maximum opportunity regret
A_1	19
A_2	13
A_3	31

The **minimax regret decision** is the one that minimises these (that is, it minimises the worse case scenario). In this case, the minimax regret decision is A_2, market the beer in moderate quantities.

	S_1	S_2	S_3
A_1	$88 - 88 = 0$	$66 - 53 = 13$	$39 - 20 = 19$
A_2	$88 - 75 = 13$	$66 - 66 = 0$	$39 - 32 = 7$
A_3	$88 - 57 = 31$	$66 - 50 = 16$	$39 - 39 = 0$

Table 2: Opportunity regret table for Duff beer example

2 Decision making with probabilities

The minimax regret strategy does not take into account any information about the probabilities of the various states of nature.

> e.g. it may be that we are almost certain to get high demand (S_1) in the Duff beer example, so A_1 might be a more appropriate action.

In many situations, it is possible to obtain probability estimates for each state of nature. These could be obtained subjectively or from empirical evidence.

Let $P(S_j)$ denote the probability that state of nature S_j occurs. Since $S_1, \ldots, S_n$ is a list of n exclusive and exhaustive events then

$$P(S_j) \geq 0 \text{ for all } j; \; \sum_{j=1}^{n} P(S_j) = 1.$$

2.1 Maximising expected monetary value

For each action A_i we can calculate the expected monetary value of the action,

$$EMV(A_i) = \sum_{j=1}^{n} \pi(A_i, S_j) P(S_j).$$

The **maximising expected monetary value strategy** involves letting

$$EMV = \max_i EMV(A_i)$$

and choosing the action A_i for which $EMV(A_i) = EMV$, that is the action which achieves the maximum expected monetary value.

2.1.1 Duff beer example

Suppose that the manager believes that the probability of strong demand (S_1) is 0.4 (so $P(S_1) = 0.4$); the probability of moderate demand (S_2) is 0.4 (so $P(S_2) = 0.4$); the probability of weak demand (S_3) is 0.2 (so $P(S_3) = 0.2$). We calculate the expected monetary value for each action.

$$\begin{aligned}
EMV(A_1) &= 88(0.4) + 53(0.4) + 20(0.2) = 60.4, \\
EMV(A_2) &= 75(0.4) + 66(0.4) + 32(0.2) = 62.8, \\
EMV(A_3) &= 57(0.4) + 50(0.4) + 39(0.2) = 50.6.
\end{aligned}$$

Thus, $EMV = \max\{60.4, 62.8, 50.6\} = 62.8$. The decision that maximises expected monetary value is A_2, to market the beer in moderate quantities.

Note that the payoff for a one-time decision will not be 62.8. It is either 75 (with probability 0.4), 66 (with probability 0.4) or 32 (with probability 0.2).

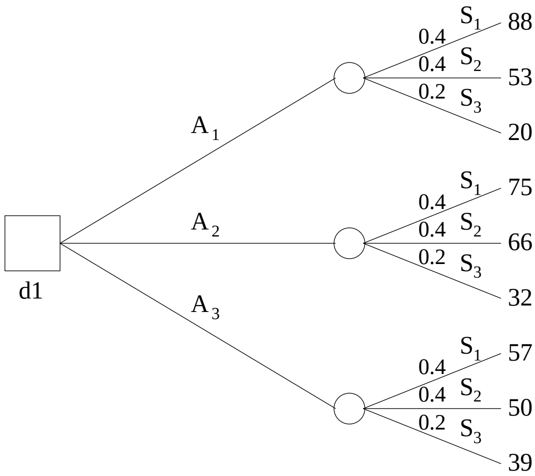

Figure 1: The decision tree for the Duff beer example.

2.2 Decision trees

A decision tree provides a graphical representation of the decision making process. It shows the logical progression that will occur over time. The tree consists of a series of nodes and branches. There are two types of node: a **decision node** (denoted as a □) which you control and a **chance node** (denoted as a ○) which you don't control. When branches leaving a node correspond to actions, then the node is a decision node; if the branches are state-of-nature branches then the node is a chance node.

2.2.1 Decision tree for the Duff beer example

The decision tree corresponding to the Duff beer example is shown in Figure 1. We solve the problem by **rollback** (or **backward dynamic programming**).

Start at the 'right' of the tree:

1. At each chance node reached, mark the node with EMV of the tree to the right of the node and remove tree to right.
2. At each decision node reached, chose decision with highest EMV, replace all of tree to right by this EMV.

Repeat until starting node reached.

The removal of subtrees is termed 'pruning'. The completed decision tree for Duff beer is shown in Figure 2

2.3 The value of perfect information

If you know precisely what state-of-nature will occur, you can maximise your payoff by always choosing the correct action. How much is this advance information worth to you? Recall that for each state S_j, $\pi^*(S_j)$ denotes the best payoff. The **expected (monetary) value under certainty (EMVUC)** is

$$EMVUC = \sum_{j=1}^{n} \pi^*(S_j)P(S_j).$$

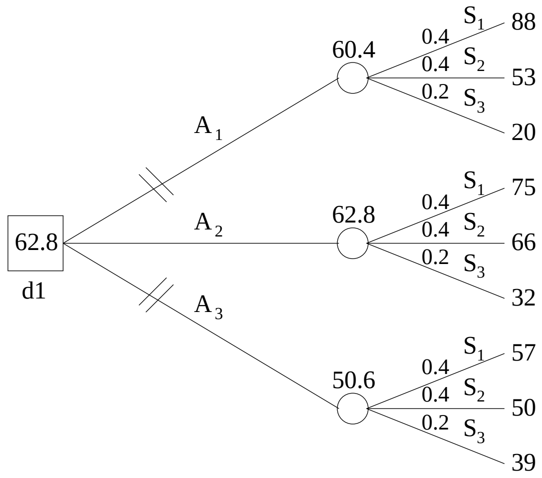

Figure 2: The completed decision tree for the Duff beer example. The non-marked branches indicate the decisions (and the ensuing consequences) we take.

The **expected value of perfect information (EVPI)**, that is the loss resulting from not having perfect information, is the difference between the expected value under certainty and the expected value under uncertainty. That is,

$$EVPI = EMVUC - EMV.$$

2.3.1 Duff beer example

Under perfect information, the manager knows before he decides whether to market the beer in large (A_1), moderate (A_2) or small (A_3) quantities if the demand is strong (S_1), moderate (S_2) or weak (S_3). His decision tree under perfect information is shown in Figure 3. Recall that $\pi^*(S_1) = 88$, $\pi^*(S_2) = 66$ and $\pi^*(S_3) = 39$, so that

$$EMVUC = 88(0.4) + 66(0.4) + 39(0.2) = 69.4.$$

Hence,

$$\begin{aligned} EVPI &= EMVUC - EMV \\ &= 69.4 - 62.8 \\ &= 6.6. \end{aligned}$$

2.4 An example with sequential decisions

Decision trees are particularly useful when the problem involves a sequence of decisions. Consider the following example. A company makes stuff and they hope for an increase in demand. In the first year, there are two possible actions: A_1 (new machinery) and A_2 (overtime). In the first year, sales can either be good (g_1) or bad (b_1) and experts judge that $P(g_1) = 0.6$, $P(b_1) = 0.4$. In the second year, the company has options which depend upon the choices made in year one. If they chose new machinery in year one, then they may choose either more machinery or more overtime. If overtime is chosen in year one, then overtime continues in year two. The sales in year two will either be high (h_2), medium (m_2)

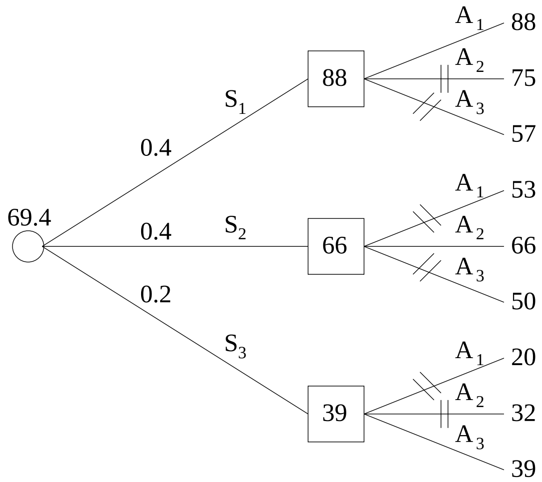

Figure 3: The completed decision tree for the Duff beer example under perfect information.

or low (l_2). If sales are good in year one, then the probability of high sales in year two is 0.5, of medium sales is 0.4 and of low sales is 0.1. If sales are bad in year one, then the probability of high sales in year two is 0.4, of medium sales is 0.4 and of low sales is 0.2. Thus, the probabilities for year two are conditional upon what happens in year one. More formally, we have:

$$P(h_2|g_1) = 0.5,\ P(m_2|g_1) = 0.4,\ P(l_2|g_1) = 0.1,$$
$$P(h_2|b_1) = 0.4,\ P(m_2|b_1) = 0.4,\ P(l_2|b_1) = 0.2.$$

The payoffs for the problem are shown in the decision tree in Figure 4. The decision tree is solved by rollback. Notice that the probabilities are conditional upon the states to the left of them are the tree (that have thus occurred).

> e.g. $EMV(A_3)$ given that we reach decision $2a$ (i.e. choose A_1 and observe g_1) is

$$EMV(A_3) \;=\; 850(0.5) + 816(0.4) + 790(0.1) \;=\; 830.4.$$

The decision rule is to choose new machinery in both years if sales in the first year are good, and new machinery in year one followed by overtime in year two if sales are bad in year one. This gives $EMV = 762.24$.

We now find $EVPI$ for the factory problem. Notice that there are six possible states for the two years $((g_1, h_2), (g_1, m_2), \ldots)$. The probabilities for these states may be found by multiplying all the probabilities on the path from the corresponding payoff to the root of the tree. For example,

$$P(g_1, h_2) \;=\; P(h_2|g_1)P(g_1) \;=\; 0.5 \times 0.6 \;=\; 0.30.$$

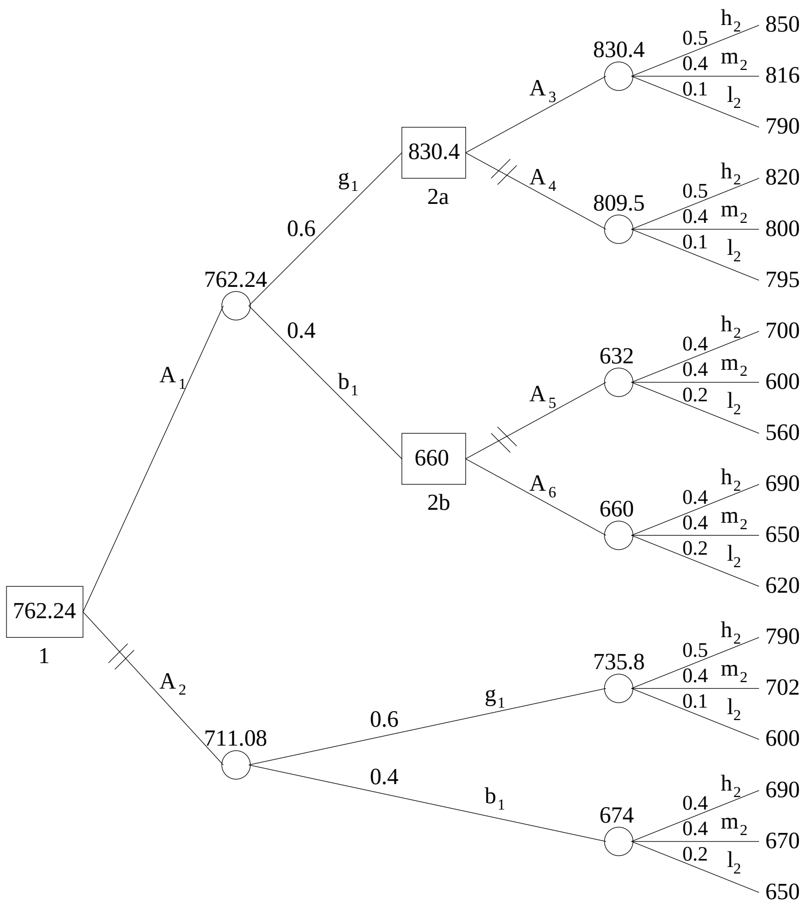

Figure 4: The completed decision tree for the factory producing stuff. A_3 and A_4 are, respectively, the actions of new machinery and overtime in year two given that we chose new machinery (A_1) in year one and observed good (g_1) sales. A_5 and A_6 are, respectively, the actions of new machinery and overtime in year two given that we chose new machinery (A_1) in year one and observed bad (b_1) sales.

We then find the best payoff for each state. For example, $\pi^*(g_1, h_1) = \max\{850, 820, 790\} = 850$. We thus have

State	Probability	Maximum Payoff
(g_1, h_2)	$0.5 \times 0.6 = 0.30$	850
(g_1, m_2)	$0.4 \times 0.6 = 0.24$	816
(g_1, l_2)	$0.1 \times 0.6 = 0.06$	795
(b_1, h_2)	$0.4 \times 0.4 = 0.16$	700
(b_1, m_2)	$0.4 \times 0.4 = 0.16$	670
(b_1, l_2)	$0.2 \times 0.4 = 0.08$	650

Hence,

$$EMVUC = (850 \times 0.30) + (816 \times 0.24) + \cdots + (650 \times 0.08) = 769.74.$$

So that,

$$EVPI = EMVUC - EMV = 769.74 - 762.24 = 7.5.$$

3 Decision analysis with sample information

It is often of value to a decision maker to obtain some further information regarding the likely state of nature. We now consider how to incorporate the sample information with the **prior** probabilities about the state of nature to obtain **posterior** probabilities. This is done via Bayes' theorem. First, we briefly review some probability theory.

3.1 A brief revision of some probability theory

An uncertain situation is one for which there are various possible outcomes and there is uncertainty as to which will occur. A sample space, denoted Ω, is a set of outcomes for the uncertain situation such that one and only one will occur. An event associated with the sample space is a set of possible outcomes, that is, a subset of the sample space. We use set notation to refer to combinations of events.

- $A \cup B$ (read as A or B) is the event that A happens or that B happens or that both A and B happen.
- $A \cap B$ (read as A and B) is the event that A happens and B happens.

A probability distribution on Ω is a collection of numbers $P(A)$ defined for each $A \subseteq \Omega$, obeying the following axioms:

1. $P(A) \geq 0$ for each event A.
2. $P(\Omega) = 1$.
3. If A and B are mutually exclusive (incompatible), that is it is impossible for A and B both to occur (i.e. if $A \cap B = \emptyset$, where $\emptyset$ denotes the empty set), then $P(A \cup B) = P(A) + P(B)$.
4. Suppose that the sample space Ω is infinite and that $A_1, A_2, \ldots$ form an infinite sequence of mutually exclusive events. Then

$$P(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} P(A_i).$$

For our purposes, the most useful consequence of these axioms is that if $A_1, A_2, \ldots, A_k$ are mutually exclusive events then

$$P(\bigcup_{i=1}^{k} A_i) \quad = \quad \sum_{i=1}^{k} P(A_i). \tag{1}$$

If $P(B) > 0$ then the conditional probability that A occurs given that B occurs is defined to be

$$P(A|B) \quad = \quad \frac{P(A \cap B)}{P(B)}. \tag{2}$$

(Intuitively, given that B occurs, it is the case that A occurs if and only if $A \cap B$ occurs. Thus, $P(A|B)$ should be proportional to $P(A \cap B)$, that is $P(A|B) = \alpha P(A \cap B)$ for some constant α. As $P(B|B) = 1$ it thus follows that $\alpha = 1/P(B)$.) From (2), we immediately deduce that $P(A \cap B) = P(A|B)P(B)$. Switching the roles of A and B in (2) and noting that $P(A \cap B) = P(B \cap A)$, gives $P(A \cap B) = P(B|A)P(A)$. Substituting this latter equation into (2) yields

$$P(A|B) \quad = \quad \frac{P(B|A)P(A)}{P(B)}. \tag{3}$$

3.2 Bayes' theorem

Let $S_1, \ldots, S_n$ be n mutually exclusive and collectively exhaustive events. Let D be some other event with $P(D) \neq 0$. We may write

$$D \; = \; \{D \cap S_1\} \cup \{D \cap S_2\} \cup \ldots \cup \{D \cap S_n\} \; = \; \bigcup_{j=1}^{n} \{D \cap S_j\},$$

where $\{D \cap S_1\}, \ldots, \{D \cap S_n\}$ are mutually exclusive. Thus, from (1) and then (2), we have

$$P(D) \; = \; \sum_{j=1}^{n} P(D \cap S_j) \; = \; \sum_{j=1}^{n} P(D|S_j)P(S_j). \tag{4}$$

(4) is often referred to as the **theorem of total probability**. From (3) and (4) we have

$$P(S_j|D) \; = \; \frac{P(D|S_j)P(S_j)}{P(D)} \; = \; \frac{P(D|S_j)P(S_j)}{\sum_{j=1}^{n} P(D|S_j)P(S_j)}. \tag{5}$$

(5) is **Bayes' theorem** and enables us to compute the posterior probabilities for each state of nature S_j following the receipt of sample information, D.

We modify $P(S_j)$, the prior probability of S_j, to $P(S_j|D)$, the posterior probability of S_j, given the sample information that D occurred, using Bayes' theorem. $P(D|S_j)$ is often termed the likelihood of S_j.

3.2.1 Duff beer example

The manager may conduct a survey that would predict strong demand (I_1), moderate demand (I_2), or weak demand (I_3). Historical data enables the manager to obtain conditional probabilities with regard to the predictions of the survey ($P(I|S)$) and these are given in

	S_1	S_2	S_3
I_1	0.8	0.3	0.3
I_2	0.1	0.5	0.1
I_3	0.1	0.2	0.6

Table 3: Conditional probabilities for survey's predictions given state of nature

Table 3. For example, $P(I_2|S_1) = 0.1$. We find the posterior probabilities of the states of nature using Bayes' theorem. Recall that our prior probabilities for the states of nature were $P(S_1) = 0.4$, $P(S_2) = 0.4$ and $P(S_3) = 0.2$. For example,

$$P(S_1|I_2) = \frac{P(I_2|S_1)P(S_1)}{P(I_2)}$$

where

$$\begin{aligned} P(I_2) &= \sum_{j=1}^{3} P(I_2|S_j)P(S_j) \\ &= (0.1 \times 0.4) + (0.5 \times 0.4) + (0.1 \times 0.2) \ = \ 0.26. \end{aligned}$$

Thus,

$$P(S_1|I_2) = \frac{(0.1 \times 0.4)}{0.26} = \frac{2}{13}.$$

It is more efficient to calculate the probabilities of the indicator, $P(I_i)$, and the posterior probabilities for the states of nature, $P(S_j|I_i)$, using a tabular approach. For each indicator I_i we form a table with the following columns:

- Column 1 counts the state, j, we are in.
- Column 2 lists the state of nature, S_j, we are in.
- Column 3 lists the prior probabilities $P(S_j)$.
- Column 4 the conditional probabilities $P(I_i|S_j)$
- Column 5 the joint probabilities $P(I_i \cap S_j)$
- Column 6 the conditional probabilities $P(S_j|I_i)$

We fill in the first four columns. Notice that the sum of all the elements in Column 3 must be 1. Since $P(I_i \cap S_j) = P(I_i|S_j)P(S_j)$, Column 5 is obtained by multiplying the corresponding row entries in Columns 3 and 4. Using the theorem of total probability, see (4), $P(I_i)$ is the sum of all of the elements in Column 5. As, see (2), $P(S_j|I_i) = P(I_i \cap S_j)/I_i)$, then Column 6 is obtained by dividing the elements in Column 5 by the sum of all the elements in Column 5. Note that the sum of all the elements in Column 6 must be 1.

We perform these for all three possible indicators, I_1, I_2, I_3, in the Duff beer example. The resulting tables are given in Table 4.

The table for I_1:

j	S_j	$P(S_j)$	$P(I_1 \vert S_j)$	$P(I_1 \cap S_j)$		$P(S_j \vert I_1)$	
1	S_1	0.4	0.8	$0.4 \times 0.8 =$	0.32	$0.32 \div 0.5 =$	0.64
2	S_2	0.4	0.3	$0.4 \times 0.3 =$	0.12	$0.12 \div 0.5 =$	0.24
3	S_3	0.2	0.3	$0.2 \times 0.3 =$	0.06	$0.06 \div 0.5 =$	0.12
		1		$P(I_1) =$	0.5		1

The table for I_2:

j	S_j	$P(S_j)$	$P(I_2 \vert S_j)$	$P(I_2 \cap S_j)$		$P(S_j \vert I_2)$	
1	S_1	0.4	0.1	$0.4 \times 0.1 =$	0.04	$0.04 \div 0.26 =$	$\frac{2}{13}$
2	S_2	0.4	0.5	$0.4 \times 0.5 =$	0.20	$0.20 \div 0.26 =$	$\frac{10}{13}$
3	S_3	0.2	0.1	$0.2 \times 0.1 =$	0.02	$0.02 \div 0.26 =$	$\frac{1}{13}$
		1		$P(I_2) =$	0.26		1

The table for I_3:

j	S_j	$P(S_j)$	$P(I_3 \vert S_j)$	$P(I_3 \cap S_j)$		$P(S_j \vert I_3)$	
1	S_1	0.4	0.1	$0.4 \times 0.1 =$	0.04	$0.04 \div 0.24 =$	$\frac{1}{6}$
2	S_2	0.4	0.2	$0.4 \times 0.2 =$	0.08	$0.08 \div 0.24 =$	$\frac{1}{3}$
3	S_3	0.2	0.6	$0.2 \times 0.6 =$	0.12	$0.12 \div 0.24 =$	$\frac{1}{2}$
		1		$P(I_3) =$	0.24		1

Table 4: Calculating the posterior probabilities of the states of nature following the predictions of the survey.

3.3 Calculating the expected payoff with sample information

We construct and solve the decision tree for the Duff beef example following the sample information. The resulting decision tree is given in Figure 5. Notice that the probabilities are conditional on all the events on a path from the probability to the root node of the tree. Thus, we require $P(S_j|I_i)$ for each i, j on the choices branches on the right side of the tree and $P(I_i)$ for each i for the choice branches coming from the root node.

If the survey predicts strong demand (I_1) then the manager should choose to market the beer in large quantities (A_2). If the survey predicts moderate (I_2) or weak demand (I_3) then the manager should choose to market the beer in moderate quantities (A_2). The expected payoff is 64.68 thousand dollars. Notice that, without the sample information, the optimal decision (in terms of maximising EMV) was to choose A_2. We can see the effect of the sample information. If the survey predicts strong demand, we should change our decision from A_2 to A_1.

For each possible piece of sample information there is a recommended action. This is called a **decision rule**. A decision rule which uses the criterion of maximising expected posterior payoff is called a **Bayes' decision rule**. We have just found the Bayes' decision rule for the Duff beer example.

3.4 The value of sample information

Obtaining sample information usually costs money. How much would we pay to receive the sample information? We should pay, at most, the expected value of sample information. Let EVwSI denote the expected value of the optimal decision *with* sample information and EV-

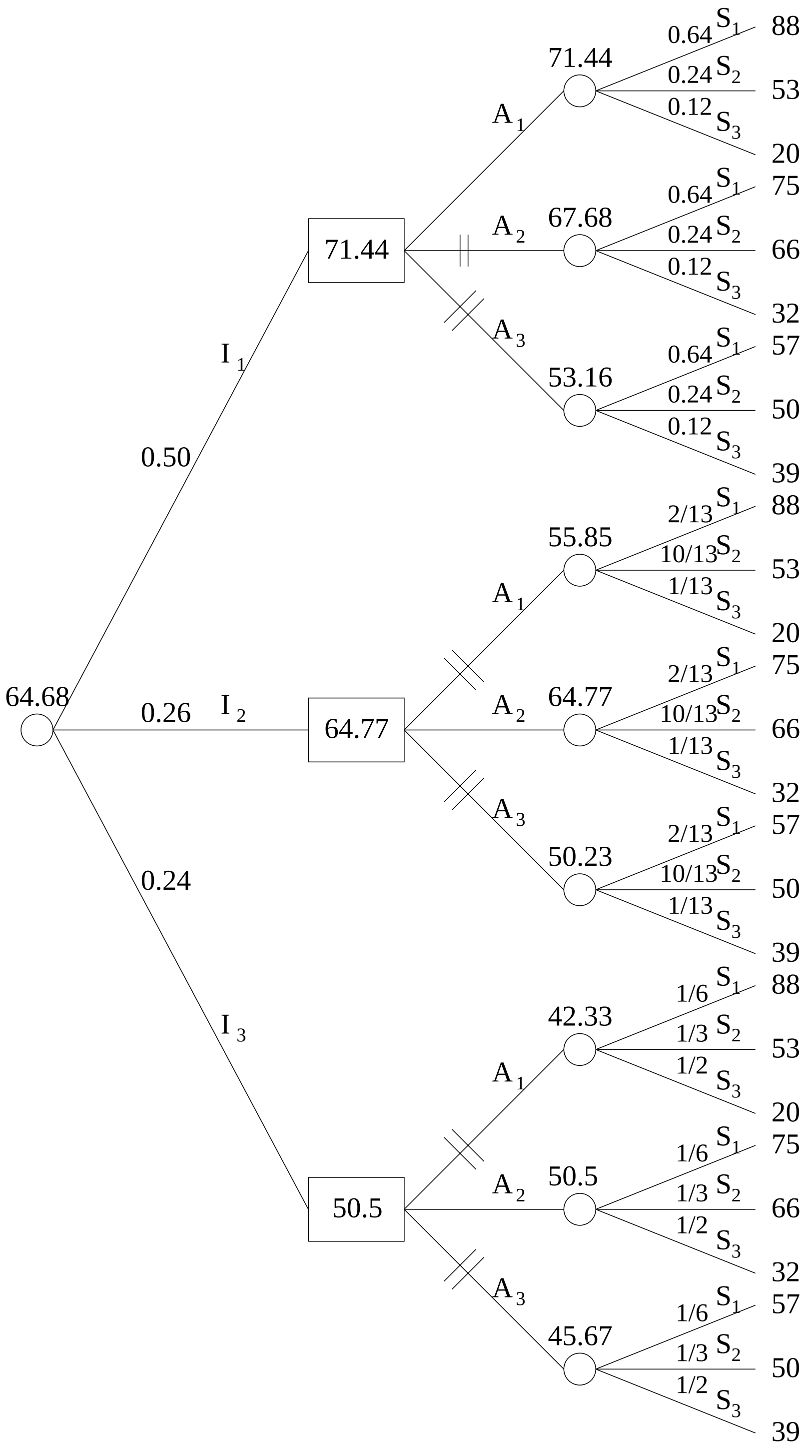

Figure 5: The decision tree for Duff beer with sample information.

woSI the expected value of the optimal decision *without* sample information. The **expected value of sample information (EVSI)** is defined as

$$EVSI \;=\; EVwSI - EVwoSI.$$

3.4.1 EVSI for the Duff beer example

In the Duff beer example, $EVwoSI = 62.8$ (see Figure 2) and $EVwSI = 64.68$ (see Figure 5). Hence,

$$EVSI \;=\; 64.68 - 62.6 \;=\; 1.88.$$

So, we would pay up to 1.88 thousand dollars for the sample information.

3.5 Efficiency of sample information

We don't expect the sample information to obtain perfect information, but we can use an efficiency measure to express the value of sample information. With perfect information having an efficiency rating of 100%, the efficiency rating E for the sample information is computed as

$$E \;=\; \frac{EVSI}{EVPI} \times 100.$$

For the Duff beer example, $EVSI = 1.88$ and $EVPI = 6.6$. The efficiency is thus

$$E \;=\; \frac{1.88}{6.6} \times 100 \;=\; 28.48\%.$$

The information from the sample is 28.48% as efficient as perfect information. Low efficiency ratings for sample information might lead you to look for other types of information before making your final decision whereas high values indicate that the sample information is almost as good as perfect information and additional sources are not worthwhile.

Suppose that there is a cost, C, for the sample information. The **net expected gain** of the sample information is the difference between between the EVSI and the cost of the sample information. The **net efficiency (NE)** is the percentage proportion of the $EVPI$ obtained by the net expected gain.

> e.g. Suppose the survey in the Duff beer example costs one thousand dollars. The net expected gain is thus
>
> $$\text{Net expected gain} \;=\; EVSI - C \;=\; 1.88 - 1 \;=\; 0.88.$$
>
> The net efficiency is
>
> $$NE \;=\; \frac{\text{Net expected gain}}{EVPI} \times 100 \;=\; \frac{0.88}{6.6} \times 100 \;=\; 13.33\%.$$

4 Utility

4.1 The St. Petersburg Paradox

So far, we have used the expected monetary value, EMV, as a means of decision making. Is this the correct approach? Consider the following game. A coin is tossed until a head appears. If the first head appears on the jth toss, then the payoff is £2^j. How much should you pay to play this game?

Let S_j denote the event that the first head appears on the jth toss. Then $P(S_j) = 1/2^j$ and the payoff for S_j, $\pi(S_j) = 2^j$. Thus,

$$EMV \;=\; \sum_{j=1}^{\infty} \pi(S_j)P(S_j) \;=\; \sum_{j=1}^{\infty} 2^j \times \frac{1}{2^j} \;=\; \infty.$$

The expected monetary payoff is infinite. However much you pay to play the game, you may expect to win more. Would you risk everything that you possess to play this game? One would suppose that real-world people would not be willing to risk an infinite amount to play this game.

The EMV criterion avoids recognising that a decision maker's attitude to the prospect of gaining or losing different amounts of money is influenced by the risks involved. A person's valuation of a risky venture is not the expected return of that venture but rather the expected utility from that venture.

4.2 Preferences

Given a choice between two rewards, R_1 and R_2 we write:

- $R_2 \prec^\star R_1$, you prefer R_1 to R_2, if you would pay an amount of money (however small) in order to swap R_2 for R_1.
- $R_1 \sim^\star R_2$, you are indifferent between R_1 and R_2, if neither $R_2 \prec^\star R_1$ or $R_1 \prec^\star R_2$ hold.
- $R_2 \preceq^\star R_1$, R_1 is at least as good as R_2, if one of $R_2 \prec^\star R_1$ or $R_1 \sim^\star R_2$ holds.

  e.g. A bakery might have five types of cake available and you assert that

  fruit cake $\prec^\star$ carrot cake $\prec^\star$ banana cake $\prec^\star$ chocolate cake $\prec^\star$ cheese cake.

  Thus, we would be willing to pay to exchange a fruit cake for a carrot cake and then pay again to exchange the carrot cake for a banana cake and so on.

We make two assumptions about our preferences over rewards. Suppose that $R_1, R_2, \ldots, R_n$ constitute a collection of n rewards. We assume

1. (COMPARABILITY) For any R_i, R_j exactly one of $R_i \prec^\star R_j$, $R_j \prec^\star R_i$, $R_i \sim^\star R_j$ holds.
2. (COHERENCE) If $R_i \prec^\star R_j$ and $R_j \prec^\star R_k$ then $R_i \prec^\star R_k$.

Comparability ensures that we can express a preference between any two rewards. Suppose that we didn't have coherence. For example, consider that carrot cake $\prec^\star$ banana cake and banana cake $\prec^\star$ chocolate cake but that chocolate cake $\prec^\star$ carrot cake. Then you would pay money to swap from carrot cake to banana cake and then from banana cake to chocolate cake. You then are willing to pay to switch the chocolate cake for a carrot cake. You are back in your original position, but have spent money to maintain the status quo. You are a **money pump**.

The consequence of these assumptions is that, for a collection of n rewards $R_1, R_2, \ldots, R_n$, there is a labelling $R_{(1)}, R_{(2)}, \ldots, R_{(n)}$ such that

$$R_{(1)} \preceq^\star R_{(2)} \preceq^\star \cdots \preceq^\star R_{(n)}.$$

This is termed a **preference ordering** for the rewards. In particular, there is a best reward, $R_{(n)}$, and a worst reward, $R_{(1)}$ (though these are not necessarily unique).

4.3 Gambles

A gamble, G, is simply a random reward.

> e.g. Toss a fair coin. If you toss a head then your reward is a carrot cake whilst if you toss a tail you receive a banana cake.

We write

$$G = p_1 R_1 +_g p_2 R_2 +_g \cdots +_g p_n R_n \tag{6}$$

for the gamble that returns R_1 with probability p_1, R_2 with probability p_2, …, R_n with probability p_n. The representation of G as in (6) is a notational convenience. We do not multiply the rewards by the probabilities nor do we add anything together. The notation $+_g$ (which you should read as 'plus gamble') helps emphasise this.

> e.g. The gamble where if you toss a head then your reward is a carrot cake whilst if you toss a tail you receive a banana cake may thus be expressed as

$$G = \frac{1}{2}\text{carrot cake} +_g \frac{1}{2}\text{banana cake}. \tag{7}$$

Note that, despite being random, each gamble can be considered a reward over which preferences are held.

> e.g. If carrot cake $\preceq^\star$ banana cake then our preferences over the gamble G given by (7) must be carrot cake $\preceq^\star G \preceq^\star$ banana cake.

We make two assumptions to ensure that our gambles are coherently compared.

1. If $R_1 \preceq^\star R_2$, $p < q$ then $pR_2 +_g (1-p)R_1 \preceq^\star qR_2 +_g (1-q)R_1$.
2. If $R_1 \preceq^\star R_2$ then $pR_1 +_g (1-p)R_3 \preceq^\star pR_2 +_g (1-p)R_3$.

Gambles provide the link between probability, preference and utility.

4.4 Utility

Preference orderings are not quantitative. For example, we can say $R_1 \prec^\star R_2$ and $R_2 \prec^\star R_3$ but not whether $R_2 \prec^\star \frac{1}{2}R_1 +_g \frac{1}{2}R_3$. For this, we need a numerical scale, your utility.

Definition 6 *(Utility)*
A utility function U on gambles $G = p_1R_1 +_g p_2R_2 +_g \cdots +_g p_nR_n$ over rewards $R_1, R_2, \ldots, R_n$ assigns a real number $U(G)$ to each G subject to the following conditions

1. *Let G_1, G_2 be any two gambles. If $G_1 \prec^\star G_2$ then $U(G_1) < U(G_2)$, and if $G_1 \sim^\star G_2$ then $U(G_1) = U(G_2)$.*

	S_1	S_2	S_3
A_1	-10,000	-10,000	-10,000
A_2	0	-100,000	-200,000

Table 5: Payoff (profit) table for the insurance example.

2. For any $p \in [0, 1]$ *and any rewards* A, B,

$$U(pA +_g (1-p)B) \quad = \quad pU(A) + (1-p)U(B).$$

REMARKS

- Condition 1. says that utilities agree with preferences, so you choose the gamble with the highest expected utility.
- Condition 2. says that, for the generic gamble $G = p_1 R_1 +_g p_2 R_2 +_g \cdots +_g R_n$, $U(G) = p_1 U(R_1) + p_2 U(R_2) + \cdots + p_n U(R_n)$. Hence, $U(G) = E(U(G))$.

  i.e. Expected utility of a gamble = Actual utility of that gamble.
- Conditions 1. and 2. combined imply that we **choose the gamble with the highest expected utility**. So, **if** we can specify a utility function over rewards, we can solve any decision problem by choosing the decision which maximises expected utility.

Note then, that if you are given utilities as a payoff rather than a monetary value (e.g. profit), to find the optimal decision: proceed as for the expected monetary value approach but with the utility values in place of the respective payoffs.

4.4.1 Example - should we insure?

The University of Bath is considering purchasing an insurance policy, against the risk of fire, for a new building. The policy has an annual cost of $£10,000$. If no insurance is purchased and minor fire damage occurs to the building then the cost is $£100,000$ while major fire damage to an uninsured building occurs a cost of $£200,000$. The probability of no fire damage (S_1) is 0.96, of minor fire damage (S_2) is 0.03 and of major fire damage (S_3) is 0.01. The action of purchasing insurance we label A_1 and that of not purchasing insurance is labelled A_2. We construct the corresponding payoff table, see Table 5.

The decision problem is a choice between two gambles over the four rewards, $£-10,000$, $£0$, $£-100,000$ and $£-200,000$. Choosing action A_1 results in a gamble, G_1, which may be expressed as

$$G_1 \;=\; 0.96£-10,000 +_g 0.03£-10,000 +_g 0.01£-10,000 \sim^\star £-10,000$$

while choosing action A_2 results in the gamble G_2,

$$G_2 \;=\; 0.96£0 +_g 0.03£-100,000 +_g 0.01£-200,000$$

The expected utility of each of these gambles is equal to the actual utility of the gamble. The respective utilities are

$$\begin{aligned} U(G_1) &= 0.96U(£-10,000) + 0.03U(£-10,000) + 0.01U(£-10,000) \\ &= U(£-10,000), \\ U(G_2) &= 0.96U(£0) + 0.03U(£-100,000) + 0.01U(£-200,000). \end{aligned}$$

We choose insurance if $U(G_1) > U(G_2)$ and we don't insure if $U(G_1) < U(G_2)$. Suppose that our utility is equal to the money value, that is

$$U(£x) \quad = \quad x.$$

The problem is then identical to maximising expected monetary value, so

$$\begin{aligned} U(G_1) &= -10,000 \;=\; EMV(A_1), \\ U(G_2) &= 0.96(0) + 0.03(-100,000) + 0.01(-200,000) \;=\; -5,000 \;=\; EMV(A_2). \end{aligned}$$

We choose action A_2. We don't purchase insurance. This is not surprising. Insurance companies make money; their premiums are always bigger than the expected payoff. Should no-one ever insure then? No, this example merely illustrates that money is not a (useful) utility scale. It doesn't capture the risk of the ventures.

Can we construct a utility function for the insurance problem such that $U(£-10,000)$, $U(£0)$, $U(£-100,000)$, $U(£-200,000)$ matches our attitude to risk? Can we construct a general utility function over a set of n rewards, $R_1, R_2, \ldots, R_n$? The answer is yes and we shall now show how.

4.5 General method for constructing utility for rewards $R_1, \ldots, R_n$

For an ordered set of rewards $R_{(1)} \preceq^\star R_{(2)} \preceq^\star \cdots \preceq^\star R_{(n)}$, define $U(R_{(1)}) = 0$ and $U(R_{(n)}) = 1$. For each integer $1 < i < n$ define $U(R_{(i)})$ to be the probability p such that

$$R_{(i)} \sim^\star (1-p)R_{(1)} +_g pR_{(n)}.$$

Thus, p_i is the probability where you are indifferent between a guaranteed reward of $R_{(i)}$ and a gamble with reward $R_{(n)}$ with probability p and $R_{(1)}$ with probability $(1-p)$. p is often termed the indifference probability for reward $R_{(i)}$. We comment that finding such a p is not easy.

4.5.1 Example revisited - should we insure?

We return to the example of §4.4.1. We have rewards

$$£-200,000 \preceq^\star £-100,000 \preceq^\star £-10,000 \preceq^\star £0.$$

over which we construct our utility function. The worst reward is $£-200,000$ and so we set $U(-£200,000) = 0$. The best reward is $£0$

We set $U(-£200,000) = 0$ and $U(-£0) = 1$. $U(-£100,000)$ is the probability p such that $-£100,000 \sim^\star (1-p)(-£200,000) +_g p(£0)$. Suppose we assert an indifference probability of 0.60. Then $U(-£100,000) = 0.6$. [Aside. Note that the expected monetary value of the gamble $(1-p)(-£200,000) +_g p(£0)$ when $p = 0.6$ is $-£80,000$. We are prepared to forfeit $£20,000$ of the EMV (i.e. $-80000-(-100000)$) to avoid risk. We are averse to risk.]

$U(-£10,000)$ is the probability q such that $-£10,000 \sim^\star (1-q)(-£200,000) +_g q(£0)$. Since $-£100,000 \prec^\star -£10,000 \prec^\star £0$ then $q \in (0.6, 1)$. Suppose we assert an indifference probability of 0.99. Then $U(-£10,000) = 0.99$.

Notice that for each action A_i and state of nature S_j we may find the corresponding utility $U(A_i, S_j)$. We can construct a utility table of our problem. This is analogous to the payoff table but with each payoff $\pi(A_i, S_j)$ replaced by the corresponding utility $U(A_i, S_j)$. The utility table for the insurance problem is given by Table 6. We choose the action which

	S_1	S_2	S_3
A_1	0.99	0.99	0.99
A_2	1	0.6	0

Table 6: Utility table for the insurance example.

maximises our expected utility.

$$\begin{aligned} EU(A_1) &= \sum_{j=1}^{3} U(A_1, S_j)P(S_j) && (8) \\ &= (0.99)(0.96) + (0.99)(0.03) + (0.99)(0.01) \\ &= 0.99; \\ EU(A_2) &= \sum_{j=1}^{3} U(A_2, S_j)P(S_j) && (9) \\ &= (1)(0.96) + (0.6)(0.03) + (0)(0.01) \\ &= 0.978. \end{aligned}$$

Thus, we choose action A_1, we insure. Compare equations (8) and (9) with (8) and (8). To calculate our expected utilities, as opposed to our expected monetary values, we simply replace each payoff by the corresponding utility.

4.6 Uniqueness of utility

We should question whether our utility is unique? Is there only one set of numbers which appropriately represent our preferences for gambles over a set of rewards? If it is not unique, then do different utilities lead to different decisions?

The answer is reassuring. A utility function on a set of rewards is unique up to an arbitrary positive linear transformation. That is

1. If $U(\cdot)$ is a utility function over a set of rewards and $a > 0$, b are numbers then $V(\cdot) = aU(\cdot) + b$ is also a utility function over that set of rewards.

2. Suppose that U, V are utilities on a reward set. Then there exist numbers $a > 0$, b such that for each reward R, $V(R) = aU(R) + b$.

Hence, if we choose the action which maximises the expected utility for U, this is also the action which maximises the expected utility for V and vice versa.