Decision Analysis

Decision making contains both psychological and rational factors. For a "satisfactory" decision, the two factors must be consistent. Decision Theory studies the rational factor in order to clarify the situation and so increase the chance of attaining consistency. The manager's ultimate success depends upon his ability to judge correctly the right blend of psychological and rational factors.

Main Elements involved in Decision Making:

(a) a number of possible **actions**, Ai, one of which has to be selected;
(b) a number of events, or **states of nature**, Sj, any one of which may hold;
(c) the value, payoff or **consequence**, Cij, to the decision maker of taking any of the available actions, in the light of the possible states of nature;
(d) the **criterion** by which the decision-maker judges between alternative actions.

		STATE OF NATURE					
		S1	S2	S3	S4	S5	S6
ACTION	A1	C11	C12	C13	C14	C15	C16
	A2	C21	C22	C23	C24	C25	C26
	A3	C31	C32	C33	C34	C35	C36

Payoff Matrix

Case 1: Decision Making under Certainty

In this case there is a single, known state of nature. Although this case appears simpler than those of non-certainty, the problem of calculating the payoff for each alternative action, or at least of identifying an action that would result in an outcome which was satisfactory, may not be trivial. Methods of Operational Research, such as Linear Programming and Dynamic Programming, may be needed.

In some cases outcomes may be characterised by several attributes which are not directly comparable. For example, when choosing between different PC's for a department, considerations such as speed, price, reliability, internal memory size, availability of software, portability and quality of graphics may all be important factors. Somehow, a set of weights for the different types of attribute has to be found, so that a single figure for the overall utility of an action can be calculated.

Case 2: Decision Making under Uncertainty

In this case there is more than one possible state of nature, but which one is the true state is not known. Furthermore, even the probabilities of the different states of nature is unknown, and cannot be estimated to any useful degree of precision.

There are several well-known approaches to decision making under uncertainty, although none is really satisfactory:

Laplace Criterion: If the probabilities of several chance events are unknown, they should be assumed equal, and the different actions should be judged according to their payoffs averaged over all the states of nature.

Maximin and Maximax Criteria: The *maximin* criterion suggests that the decision-maker should choose the alternative which maximises the minimum payoff he can get. This pessimistic approach implies that the decision-maker should expect the worst to happen.

The *maximax* criterion indicates that the decision-maker should, on the contrary, choose the alternative which maximises the maximum value of the outcome. This optimistic approach implies that the decision-maker should assume the best of all possible worlds.

Since either the maximin or maximax approaches alone focus too narrowly on a single element in what may be a large payoff matrix, it has been proposed that the two criteria be combined. The decision-maker could take into consideration both the largest and smallest payoffs and then weigh their importance according to an *index of optimism*, between 0 and 1.

Regret Criterion: The *regret* of an outcome is the difference between the value of that outcome and the maximum value of all the possible outcomes, in the light of the particular chance event that actually occurred. The decision-maker should choose the alternative that minimises the maximum regret he could suffer.

<u>Example</u>: A computer company is planning to embark on a major TV advertising campaign. Three TV networks are available for use by the firm in carrying out its campaign. These alternative networks are denoted by A1, A2 and A3. Associated with each TV network are three possible outcomes. These represent increments in total profits for the firm, which result from the use of a particular TV network. The payoff matrix is:

payoff (in £m)	S1	S2	S3
A1	35	65	5
A2	36	12	48
A3	-3	60	30

Laplace Criterion: Assuming each outcome is equally likely, the average monetary payoffs of A1, A2 and A3 are 35, 32 and 29, respectively. Accordingly, the firm should select A1.

Maximin and Maximax Criteria: The minimum payoffs are 5, 12 and -3, and so, by the maximin criterion, the firm should select A2. The maximum payoffs are 65, 48 and 60, and so, according to the maximax criterion, the firm should select A1. With a coefficient of optimism of 0.6, the weighted minima and maxima are 41 (= 0.6 x 65 + 0.4 x 5), 33.6 (= 0.6 x 48 + 0.4 x 12) and 34.8 (= 0.6 x 60 - 0.4 x 3), and so the firm should select A1.

Regret Criterion: The regret matrix is

regret (in £m)	S1	S2	S3
A1	1	0	43
A2	0	53	0
A3	39	5	18

The maximum regret for A1, A2 and A3 is 43, 53 and 39, respectively. Under this criterion, A3 should be chosen.

Weaknesses of the Criteria for Choosing under Uncertainty

Suppose that it is required to specify the nature of an emergency air cargo fleet for supplying relief to the victims of earthquakes. Assume that two kinds of aircraft are available, differing only in that one has a longer-range than the other. If a crisis should develop relatively close to home base (London, for example), then the short-range plane would be most effective. If the crisis is far away, however, then the short-range plane might be forced to take an indirect route and thus be inefficient. Finally, someone suggests that the required emergency carrying capacity could be attained economically, for some situations at least, by using trucks instead of aircraft.

For an example set of possible emergencies, assume that the payoff matrix is:

	Iberian Peninsula	Azerbaijan	Wales
Short haul	100	40	30
Long haul	70	80	20
Trucks	0	0	110

where the numbers in the matrix represent some acceptable measure of utility.

Using the maximin decision rule, we calculate the minimum payoff for each choice as 30, 20 and 0, and then select short haul aircraft as the best choice because it maximises (at 30) the minimum value of the outcome.

Using the maximax decision rule, we calculate the maximum payoff for each choice as 100, 80 and 110, and then select trucks as the best choice because its maximum value of 110 provides the best return if everything goes right.

Both the maximin and maximax criteria imply extreme attitudes to risk. Maximin assumes that the decision-maker is very cautious, whilst maximax presumes that the decision maker is almost foolhardy.

Combining maximin and maximax with an index of optimism 0.6, we calculate:

	Optimism		Pessimism		Total
Short haul	0.6 x 100	+	0.4 x 30	=	72
Long haul	0.6 x 80	+	0.4 x 20	=	56
Trucks	0.6 x 110	+	0.4 x 0	=	66

Short haul aircraft is again the recommended option. Although this combined criterion seems better than the pure maximin or maximax approaches, it still discards all information about outcomes with intermediate values, although these may be the most interesting ones.

The regret matrix is:

	Iberian Peninsula	Azerbaijan	Wales
Short haul	0	40	80
Long haul	30	0	90
Trucks	100	80	0

The regret criterion calculates the maximum regret for the three actions as 80, 90 and 100, and leads to the choice of short haul aircraft, as this minimises (at 80) the maximum regret.

The serious flaw with the regret criterion is that the values of the regrets are not absolute. They are strictly relative to other alternatives and will vary as the number of system alternatives is expanded or contracted.

The choice between serious alternatives can, in fact, be altered by introducing irrelevant or flippant choices. Such is the case in the example problem. Suppose, as seems reasonable, that the alternative of trucks is irrelevant to the question of what aircraft belong in an emergency fleet. Dropping the trucks alternative from consideration changes the regret matrix to:

	Iberian Peninsula	Azerbaijan	Wales
Short haul	0	40	0
Long haul	30	0	10

and the choice of aircraft changes from short haul to long haul.

The unreasonable effect of irrelevant alternatives is not easily corrected since the dependence of the value of regret is a function of the alternatives compared, and it may be impossible to obtain a consistent, transitive ranking of alternatives. Even using pairwise comparisons, in this case, long haul is preferred to short haul, short haul is preferred to trucks, but trucks are preferred to long haul.

	Iberian Peninsula	Azerbaijan	Wales
Short haul	0	0	80
Trucks	100	40	0

	Iberian Peninsula	Azerbaijan	Wales
Long haul	0	0	90
Trucks	70	80	0

This kind of intransitivity alone is sufficient reason to be very sceptical about the use of the regret criterion for decision making.

Under the Laplace criterion, every chance event (or state of nature) is assumed equally likely. The expected values of the three actions are:

E(short haul)	=	l/3(100 + 40 + 30)	= 56.67
E(long haul)	=	l/3(70 + 80 + 20)	= 56.67
E(trucks)	=	l/3(0+0+ 110)	= 36.67

Therefore, short haul and long haul are equally recommended under this criterion.

The peculiar flaw of the Laplace criterion is that it is sensitive to the description of chance events, and can be altered by the introduction of irrelevant or trivial possibilities. Consider the reduced payoff matrix:

	Iberian Peninsula	Azerbaijan
Short haul	100	40
Long haul	70	80

Taking each chance event as equally likely, as indicated by the Laplace criterion, the expected values for short haul and long haul are 70 and 75, so that long haul appears better. Now split the Iberian Peninsula into Spain and Portugal:

	Spain	Portugal	Azerbaijan
Short haul	100	100	40
Long haul	70	70	80

The expected values for short haul and long haul are now 80 and 73.33, and so short haul is preferred.

The introduction of this alternative way of categorising the states of nature has essentially, but covertly, had the effect of changing the probability of a nearby disaster from l/2 to 2/3. Would it not be preferable to make such judgements explicitly, rather than leave them to an unthinking mechanism?

Case 3: Decision Making under Risk

Risk refers to the situation in which the outcome of each action is not certain, but where the probabilities of the different states of nature (and hence of the alternative outcomes) can be determined.

Introduction to Probability

The **sample space**, S, for a situation or experiment is the set of all possible basic outcomes. For example, if an ordinary die is thrown once then

$$S = \{1, 2, 3, 4, 5, 6\}$$

which is the set of all the possible numbers that could be thrown.

An **event** is any set of possible basic outcomes. For example:

A: "Throwing an even number" = {2, 4, 6}
B: "Throwing a number greater than 4" = {5, 6}

are examples of events.

A set of events is **mutually exclusive** if no two can occur at the same time. For example, {1, 3}, {2, 4} and {6} are three mutually exclusive events.

A set of events is **exhaustive** if at least one of them is bound to occur. For example, {1, 2, 3}, {3, 4, 5} and {5, 6} are three exhaustive events.

The **probability** P(E) of an event E is an indication of how likely that event is to happen.

Probabilities have the following properties:

- For any event E: $0 \le \mathrm{P}(E) \le 1$.
  (i.e., nothing can have less than 0% chance or more than 100% chance of occurring).
- If an event is impossible then its probability is 0 .
- If an event is certain then its probability is 1 .
- If A and B are mutually exclusive events then: $\mathrm{P}(A \text{ or } B) = \mathrm{P}(A) + \mathrm{P}(B)$,
  (where "A or B" is the event that one or the other occurs).
- The sum of the probabilities of a mutually exclusive and exhaustive set of events is 1 .

If A and B are events, then the event "A and B" (that both occur) is called the **joint event**.

The **conditional probability** of event A, given that event B occurs, P(A | B) , is defined by:

$$\mathrm{P}(A \mid B) = \frac{\mathrm{P}(A \text{ and } B)}{\mathrm{P}(B)}$$

It follows that the joint probability $\mathrm{P}(A \text{ and } B) = \mathrm{P}(A \mid B) \times \mathrm{P}(B)$

If B is any event and $\{A1, A2, \dots , An\}$ are mutually exclusive and exhaustive events, then we have the **generalized addition law**:

$$\begin{aligned}\mathrm{P}(B) &= \mathrm{P}(B \text{ and } A1) + \mathrm{P}(B \text{ and } A2) + \dots + \mathrm{P}(B \text{ and } An) \\ &= \mathrm{P}(B \mid A1) \times \mathrm{P}(A1) + \mathrm{P}(B \mid A2) \times \mathrm{P}(A2) + \dots + \mathrm{P}(B \mid An) \times \mathrm{P}(An)\end{aligned}$$

Two events are said to be **independent** if the occurrence of one has no effect on the probability of the occurrence of the other.

If events *A* and *B* are independent, then: $P(A \mid B) = P(A)$ and so

$$P(A \text{ and } B) = P(A) \times P(B)$$

Three are three **alternative interpretations of probability**:

Classical (or theoretical) probability:

Based on simple games of chance involving symmetric objects such as fair coins, dice and packs of cards, in which basic outcomes are equally likely. For example, for a fair die:

$$P(1) = P(2) = \ldots = P(6) = 1/6 .$$

In such simple cases, with finite sample spaces (i.e., with a finite number of basic outcomes), the probability of an event *E* is simply:

$$P(E) = \frac{\text{number of basic outcomes in } E}{\text{number of basic outcomes in sample space}}$$

For example, if E is the event of throwing a number greater than 4 with a fair die, then

$$P(E) = 2/6 = 1/3 .$$

Long-term frequency:

Based on observing *n* repeated trials of an experiment and counting the number of times *m* that a particular event *E* occurs. The relative frequency, m/n , with which the event occurs is an estimate of the probability of *E*. As *n* is increased, this ratio becomes a more and more accurate estimate of the probability. This interpretation applies to a far wider range of phenomena than does classical probability, for example to industrial processes and to life insurance.

Subjective probability:

A measure of the degree of belief one has that an event will occur. It is a personal judgment, based on all relevant information one has at the time, such as a bookie's odds in a horse race. This interpretation is applicable to the widest range of phenomena as it neither requires symmetry (as in classical probability) nor repeatability of identical trials (as in the frequentist approach). It is the most appropriate interpretation in the area of management decision making.

Bayes' Theorem

If *A* and *B* are any two events then: $P(A \text{ and } B) = P(A \mid B) \times P(B)$

Similarly: $P(B \text{ and } A) = P(B \mid A) \times P(A)$

But the joint event "*A* and *B*" is the same as "*B* and *A*" and so we can equate the probabilities.
Therefore:

$$P(A \mid B) \times P(B) = P(B \mid A) \times P(A)$$

and so:

$$P(A \mid B) = \frac{P(B \mid A)}{P(B)} \times P(A)$$

Bayes' theorem can be seen as relating $P(A \mid B)$, the conditional probability of *A* given *B*, to the absolute probability P(*A*).

If {*A1*, *A2*, … , *An*} are mutually exclusive and exhaustive, then for any one of the *Ai* :

$$P(Ai \mid B) = \frac{P(B \mid Ai) \times P(Ai)}{P(B \mid A1) \times P(A1) + \ldots + P(B \mid An) \times P(An)}$$

We often use Bayes' Theorem when we want to find the probability of a particular state of affairs, in the light of observations or experiments that have been made. If *A1*, *A2*, … , *An* are alternative states and observation *B* has been made, the last equation shows how to relate the conditional probability that state *Ai* is the true state, given observation *B*, to the "absolute" probability of *Ai*, estimated before *B* had been observed.

P(*Ai*) is sometimes called the **prior probability** of *Ai* - prior to collecting any extra relevant information or making any extra observations related to the possible occurrence of *Ai*.

P(*Ai* | *B*) is sometimes called the **posterior probability** of *Ai* - posterior (i.e., after) observing that *B* occurred.

P(*B* | *Ai*) is sometimes called the **likelihood** of *B* - the probability of observing *B*, in state *Ai*.

So Bayes' Theorem provides the mechanism for updating our estimate as to the chance of *Ai* being the true state, in the light of new information.

Example: Two opaque jars *X* and *Y* each contain ten balls. Jar *X* contains four red balls and six green balls. Jar *Y* contains eight red balls and two green balls. One of the two jars is chosen at random. What is the probability that jar *X* was chosen?

In the absence of any further information, the answer is 1/2. The prior probability:

$$P(X) = 1/2.$$

Now suppose that some information is collected to help in deciding which jar was chosen. A ball is taken at random from the jar. It turns out to be a red ball. What effect does this have on the assessment of the probability that jar *X* had been chosen?

If jar *X* had been chosen, the probability that a red ball would be withdrawn is 4/10. If jar *Y* had been chosen, the probability that a red ball would be withdrawn is 8/10. The likelihoods of a red ball being withdrawn are:

$$P(red \mid X) = 4/10, \quad P(red \mid Y) = 8/10.$$

We are interested in the posterior probability P(*X* | *red*). Using Bayes' Theorem:

$$\begin{aligned} P(X \mid red) &= \frac{P(red \mid X) \times P(X)}{P(red \mid X) \times P(X) + P(red \mid Y) \times P(Y)} \\ &= \frac{4/10 \times 1/2}{(4/10 \times 1/2) + (8/10 \times 1/2)} \\ &= \frac{1/5}{1/5 + 2/5} \\ &= 1/3 \end{aligned}$$

It is sometimes easier to employ a diagrammatic method based on Bayes' theorem, which involves drawing a unit square, marking the a priori probabilities along the base, and dividing the square into corresponding vertical rectangles. We then divide each rectangle according to the likelihood values, and calculate the area of each of the resulting smaller rectangles.

The two original vertical rectangles correspond to the two events: "jar *X* was picked" and "jar *Y* was picked". Each of the smaller rectangles represents a joint event, such as "jar *X* was picked **and** a *red* ball was drawn from it". The areas of the rectangles are the probabilities of these different events.

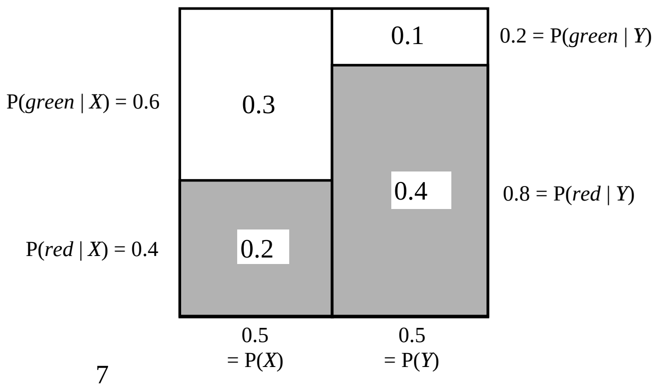

The shaded rectangles represent the two joint events which involve drawing a red ball. The overall probability of such a result is:

$$\begin{aligned} P(red) &= P(X \text{ and } red) + P(Y \text{ and } red) \\ &= P(red \mid X) \text{ x } P(X) + P(red \mid Y) \text{ x } P(Y) \\ &= 0.4 \text{ x } 0.5 + 0.8 \text{ x } 0.5 \\ &= 0.2 + 0.4 \\ &= 0.6 \end{aligned}$$

The unshaded regions represent the joint events that involve drawing a green ball:

$$\begin{aligned} P(green) &= P(X \text{ and } green) + P(Y \text{ and } green) \\ &= P(green \mid X) \text{ x } P(X) + P(green \mid Y) \text{ x } P(Y) \\ &= 0.6 \text{ x } 0.5 + 0.2 \text{ x } 0.5 \\ &= 0.3 + 0.1 \\ &= 0.4 \end{aligned}$$

The posterior probabilities are:

$$P(X \mid red) = \frac{P(X \text{ and } red)}{P(red)} = \frac{P(red \mid X) \times P(X)}{P(red)} = \frac{0.2}{0.6} = 1/3$$

$$P(Y \mid red) = \frac{P(Y \text{ and } red)}{P(red)} = \frac{P(red \mid Y) \times P(Y)}{P(red)} = \frac{0.4}{0.6} = 2/3$$

$$P(X \mid green) = \frac{P(X \text{ and } green)}{P(green)} = \frac{P(green \mid X) \times P(X)}{P(green)} = \frac{0.3}{0.4} = 3/4$$

$$P(Y \mid green) = \frac{P(Y \text{ and } green)}{P(green)} = \frac{P(green \mid Y) \times P(Y)}{P(green)} = \frac{0.1}{0.4} = 1/4$$

Thus, drawing a red ball decreases the probability of jar X from 1/2 to 1/3 and increases the probability of jar Y from 1/2 to 2/3. Drawing a green ball increases the probability of jar X from 1/2 to 3/4 and decreases the probability of jar Y from 1/2 to 1/4.

The knowledge of the probabilities permits the calculation of the expected values for the alternatives, and thus a rational selection between them.

Expected Monetary Value (EMV) Criterion: For each action in turn, the expected payoff in cash terms is calculated. The action with the highest expected payoff is the preferred choice. To derive the expected payoff of an action, each conditional value is multiplies by its probability of occurrence, and the sum of all such products is taken. In symbols, the expected payoff of an action Ai is:

$$E_i = \sum_j C_{ij} \text{ x } P(S_j)$$

where $P(S_j)$ is the probability of state of nature S_j.

<u>Example</u>: A computer company must decide whether or not to undertake some research work, which would entail a capital expenditure of £800,000. The two possible states of nature are that the research succeeds, or that it fails. If it succeeds, it would lead to a cash return of £2m. However, unsuccessful research would not yield any return. It is estimated that there is a 0.6 probability of the research being successful.

If research were conducted which succeeded, then the net payoff would be £2m - £0.8m = £1.2m. If research were conducted which failed, then there would be a net loss of £0.8m. On the other hand, if research were not conducted, then there would be a zero payoff, no matter whether the research would have been successful or not.

payoff (in £m)	Research Succeeds	Research Fails	Expected Payoff
Conduct Research	1.2	-0.8	0.4
Don't do Research	0	0	0
Probability	0.6	0.4	

The expected payoff from conducting research is: £1.2m x 0.6 + (-£0.8m) x 0.4 = £0.4m. The expected payoff from not conducting research is zero. Hence, under the EMV criterion, the former action is to be preferred.

Decision Trees: Decision trees are a useful tool in representing *multi-stage decision problems*. Possible actions at any point are shown as branches emanating from a *decision point*, represented by a small square. The various possible outcomes of an action are shown as branches leading from a *chance point*, designated by a node with a circle, at the end of the branch for the action. Probabilities are associated with each branch from a chance point to an outcome. Values associated with each outcome are shown at the ends of the corresponding branches.

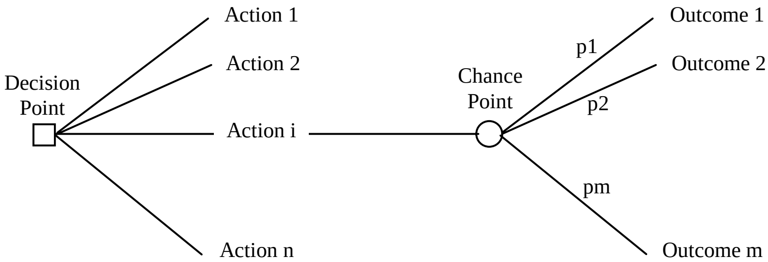

Having constructed the tree, one can identify the course of action prescribed by the EMV criterion, using the method of *folding back and pruning*. Folding back consists of calculating the EMV for each node. We start with the nodes which have no successors. These are the ones which are the furthest in the future. A decision node with no successor chance node would normally be one for which there is no uncertainty as to which action to choose. The value of the decision node is taken to be the payoff of that action. For a chance node with no successor decision node, we assign to it the expected value of the payoffs associated with the branches emanating from that chance node.

For a decision node all of whose successor chance nodes have been evaluated, we choose the action which leads to a chance node with the highest payoff. All the other possible actions for that decision node are removed from consideration and their branches pruned. The decision node is given the value of the chance node at the end of the chosen action. For a chance node all of whose successor decision nodes have been evaluated, we calculate the expectation of all the successor decision nodes' values.

This process is repeated systematically until we have evaluated the decision node at the root of the tree (which is the decision we have to make immediately). Those parts of the decision tree that can be reached from the root via unpruned branches provide a complete solution to the multi-decision problem.

Example: A company must decide whether to build a small plant or a large one, to manufacture a new product with a market life of ten years.

Demand for the product may possibly be high during the first two years but, if many of the initial users find it unsatisfactory, the demand could then fall to a low level thereafter. Alternatively, high initial demand could indicate the possibility of a sustained high-volume market. If the demand is initially high and remains so, and the company finds itself with insufficient capacity within the first two years, competing products will certainly be introduced by other companies.

If the company initially builds a big plant, it must live with it for the whole ten years, whatever the size of the market demand. If it builds a small plant, there is the option of expanding the plant in two years time, an option that it would only take up if demand were high during the introductory period. If a small plant is built initially and demand is low during the introductory period, the company will maintain operations in the small plant, and make a good profit on the low-volume throughput.

Marketing Information: 60% chance of a large market in the long run, and 40% of a long-term low demand developing initially as follows:

Initially High, sustained High	60%	
Initially High, long-term Low	10%	long-term Low 40%
Initially Low, continuing Low	30%	
Initially Low, long-term High	0%	

Annual Income:

(a) A large plant with high market volume would yield £1m annually for ten years.
(b) A large plant with low market volume would yield only £0.1m annually.
(c) A small plant with low market demand would yield a cash income of £0.4m per annum.
(d) A small plant during an initial period of high demand would yield £0.45m per annum but, because of competition, this would drop to £0.25m per annum in the long run, if high demand continued.
(e) If an initial small plant were expanded after two years to meet sustained high demand, it would yield £0.7m annually for the remaining eight years.
(f) If an initial small plant were expanded after two years, but high demand were not sustained, the estimated annual income for the remaining eight years would be £0.05m.

Capital Costs: A large plant would cost £3m to build. A small plant would cost £1.3m initially, and an additional £2.2m if expanded after two years.

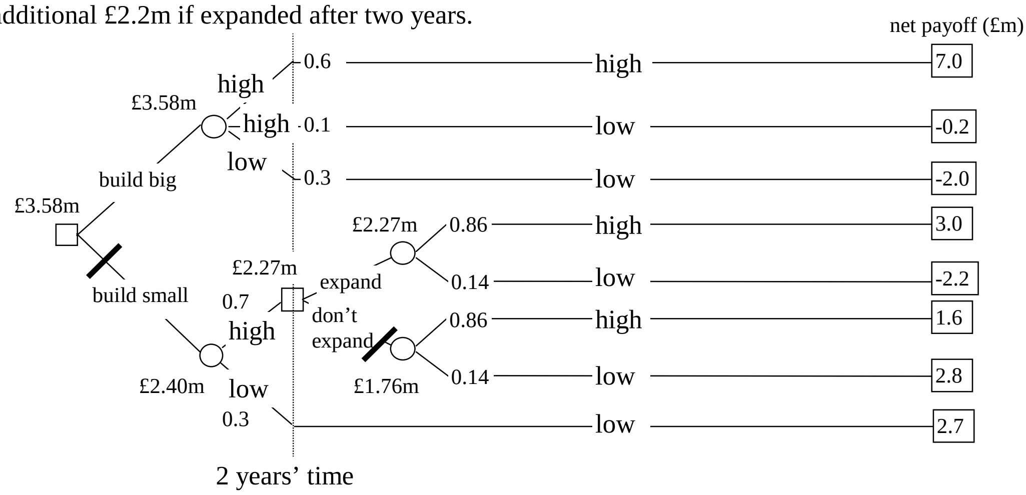

Since a state of nature corresponding to an initially low market volume, followed by a long-term high one has probability zero, branches corresponding to this state of nature have been omitted from the above decision tree.

The net payoff figure of £7m for the case of ten years' high market volume is obtained by subtracting the capital cost of £3m from the total income of £10m. Net payoff figures corresponding to the other possible outcomes are calculated in a similar fashion.

The top branch emanating from the chance node at the end of the initial "build big" action corresponds to a state of nature in which market volume is high for all ten years. The probability for this branch is taken directly from the marketing information. The probabilities for the other two branches coming from this chance node (*initially high, long-term low* and *initially low, long-term low*) are obtained similarly.

The branches coming from the chance node at the end of the initial "build small" action correspond to initially high and initially low market volumes. Again, the probabilities can be obtained from the marketing information.

However, the probabilities for the chance nodes following the contingent “expand” and “don’t expand” actions are a little more involved. Consider the top branch emanating from the chance node following “expand”. This corresponds to a state of nature of a high long-term market volume, given that there was a high volume in the first two years (and that the company expanded their plant). The probability of this state of nature is:

P(long-term high | initially high) = P(initially high & sustained high) / P(initially high)
= 0.6 / 0.7
= 0.86

Similarly, the lower branch emanating from the chance node following “expand” corresponds to low long-term market volume, given that there was a high volume in the first two years. This has probability:

P(long-term low | initially high) = P(initially high & long-term low) / P(initially high)
= 0.1 / 0.7
= 0.14

The branches emanating from the chance node following “don’t expand” have similar probabilities.

The chance node following “build big” has expected value:
7.0 x 0.6 + (-0.2) x 0.1 + (-2.0) x 0.3 = 4.2 - 0.02 - 0.6 = £3.58m

Similarly, the chance nodes following “expand” and “don’t expand” have values £2.27m and £1.76m.

Considering the decision point at the end of the second year, in the circumstance when the company built small initially, but the initial market volume was high, there are two possible actions: “expand” and “don’t expand”. These two actions lead to chance nodes with expected values £2.27m and £1.76m, respectively. An expected payoff of £2.27m is to be preferred to one of £1.76m, and so “expand” is chosen in preference to “don’t expand”. What this means is that if the company builds small initially, and the initial volume is high, then (according to the EMV criterion) the company should expand their plant at the end of two years. The “don’t expand” branch is pruned, and the decision node is given the expected value £2.27m.

The chance node following “build small” can now be evaluated. Two branches emanate from this node, corresponding to the two states of nature “initially high volume” and “initially low volume”. The high branch, with probability 0.7, leads to the decision node to which we have just assigned the value £2.27m. The low branch, with probability 0.3, has an outcome with a payoff £2.7m. The expected value for this chance node is: 2.27 x 0.7 + 2.7 x 0.3 = 1.59 + 0.81 = £2.4m.

We finally consider the decision node at the root of the tree. There are two possible actions “build big” and “build small”. The former leads to a chance node with an expected value £3.58m. The latter to the chance node discussed in the previous paragraph, which has the expected value £2.4m. The action “build big” leads to the greater expected monetary value, and so this is the preferred action. The “build small” action is pruned, and the decision node is given the expected value £3.58m.

Hence the recommended course of action is for the company to build a big plant right from the start. The expected net payoff if they do so is £3.58m.

Using Bayes’ Theorem: All probabilities are essentially relative, i.e., their values depend on the information possessed by the decision maker when assessing their values.

The values of probabilities may change as new information reaches the decision maker. Hence the expected payoffs of various acts may change, and so there may be a different optimal act. The mechanism by which probabilities change, in the light of new information, can be studied using Bayes’ Theorem.

The Value of Information: The purpose of information is to throw light on a situation, helping the decision maker to choose a sensible action on the basis of more reliable estimates of the probabilities involved, and reducing the risk of making an expensive mistake. Although it is a straightforward matter of accounting to calculate the cost of a piece of information, the value of that information is subjective and dependent on the decision making situation.

Example of Venture Analysis: A company must decide whether or not to launch a particular new product. They think that there is a 70% chance that the demand for the product will be high, and a 30% chance that it will be low. If the product is launched and has a high demand, the net profit will be £500,000. If the product is launched and the demand is low, there will be a net loss of £250,000.

The company could arrange for market research to be carried out before making a final decision about launching the product. In the past, similar research has been 85% successful in correctly forecasting a high market share, and 75% successful in correctly forecasting a low market share. The cost of research is £10,000.

What should the company do?

The payoff matrix is:

profit in £000s	High market share	Low market share	Expected Profit
Launch Product	500	-250	275
Drop Product	0	0	0
Probability	0.7	0.3	

The expected monetary value (EMV) of launching the product is: 500 x 0.7 + (-250) x 0.3 = £275,000. The EMV of dropping the product is: 0 x 0.7 + 0 x 0.3 = £0. Hence, in the absence of any further information, the preferred act (according to the EMV criterion) would be to launch the product, and this would lead to an expected profit of £275,000.

If we were certain that the outcome would be “high market share” then we would launch the product, leading to a profit of £500,000, since such a profit is preferred to a profit of £0. If we were certain that the outcome would be “low market share” then we would drop the product, leading to a profit of £0, since breaking even is preferred to a loss of £250,000 (i.e., a profit of -£250,000).

Suppose that, somehow or other, we could obtain 100% reliable information about the nature of the outcome. There is an a priori belief of 0.7 that the outcome will be high and a belief of 0.3 that the outcome will be low. Thus there is a probability of 0.7 that perfect information would tell us that the outcome would be high, and a probability of 0.3 that perfect information would tell us that the outcome would be low.

If we possessed perfect information and took the best action, depending on what the information told us, there would be a 0.7 probability of obtaining a profit of £500,000 and a 0.3 probability of breaking even. Hence, the expected profit when in possession of perfect information is: 500 x 0.7 + 0 x 0.3 = £350,000.

Therefore, the ***Expected Monetary Value of Perfect Information*** (EMVPI) is the expected payoff when in possession of perfect information minus the best expected payoff we could obtain without any extra information: 350 - 275 = £75,000.

This figure of £75,000 represents the value of information that is 100% reliable. In no circumstances would it ever be worthwhile obtaining information that cost more than this.

The proposed market research is not 100% reliable. We shall refer to a research result that forecasts a high market share as *optimistic* and one that forecasts a low market share as *pessimistic*. According to the record of previous research undertaken:

P(optimistic | high) = 0.85, P(pessimistic | high) = 0.15,
P(optimistic | low) = 0.25, P(pessimistic | low) = 0.75.

These likelihood figures can be combined with the a priori probabilities of a high or low market share to obtain a posteriori probabilities:

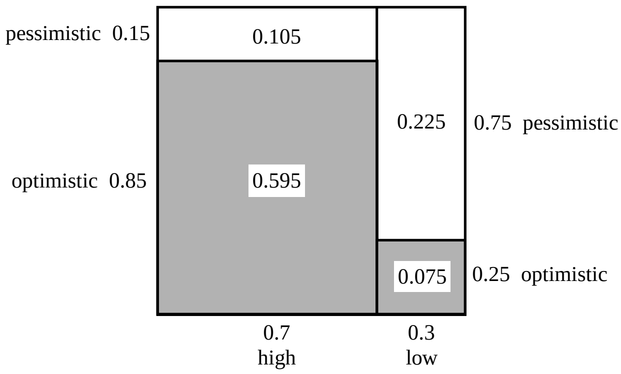

P(optimistic) = 0.595 + 0.075 = 0.67,
P(pessimistic) = 0.105 + 0.225 = 0.33

P(high | optimistic) = 0.595/0.67 = 0.89, P(low | optimistic) = 0.075/0.67 = 0.11
P(high | pessimistic) = 0.105/0.33 = 0.32, P(low | pessimistic) = 0.225/0.33 = 0.68

In order to find the expected value of the information from marketing research, we use a decision tree representing the different actions and states of nature. In calculating the payoffs, we ignore (for the moment) the cost of the marketing research.

From the tree, we see that, if we first commission the research and act rationally after receiving its results, the expected profit will be £278,500 not counting the cost of the research itself. On the other hand, the best that we can do without any further information is to launch the product and expect a profit of £275,000.

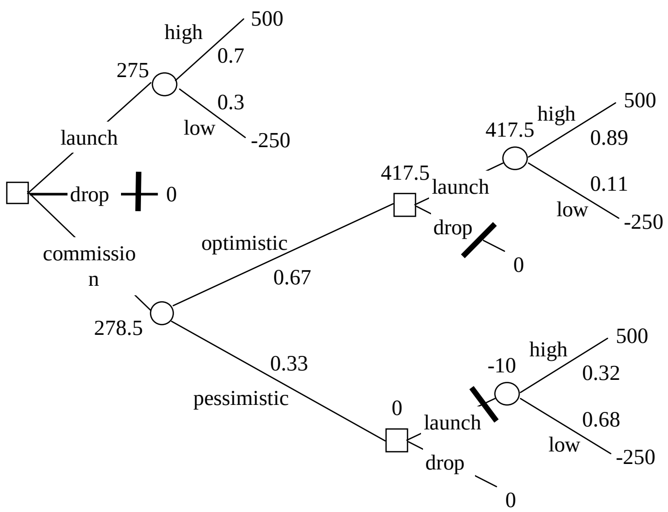

Hence, the ***Expected Monetary Value of Sample Information*** (EMVSI) is: 278.5 - 275 = £ 3,500. This compares to a *cost* of £10,000. Hence, the research should not be commissioned, but the product launched straight away.