Ch5. Decision Theory

In this chapter we will introduce the concept of **Decision Theory** briefly.

5.1. Expectation (mean)

- If X is a discrete random variable with the probability function

$$P(X = x_i) = p_i, \text{ for } i = 1, 2, \ldots,$$

the **expected value** (or **expectation**, or **mean**) of X is

$$E(X) = \sum_i x_i p_i.$$

5.1 Expectation

- If X is a continuous random variable with density function $f(x)$, the **expected value** of X is

$$E(X) = \int_{-\infty}^{\infty} xf(x)dx.$$

Properties

- If a and b are constant, then $E(aX + b) = aE(X) + b$;
- $E(X + Y) = E(X) + E(Y)$.

5.1 Expectation

Example 1. If X has a Binomial distribution $B(n, p)$, then

$$p_i = P(X = i) = \binom{n}{i} p^i (1-p)^{n-i},$$

then

$$E(X) = \sum_{i=1}^{n} i p_i = \sum_{i=1}^{n} \left[i \frac{n!}{(n-i)!i!} p^i (1-p)^{n-i} \right] = \ldots = np.$$

5.1 Expectation—Example: Roulette, Bets and Odds

The outcome of Roulette is

- one of number in 0 to 36;
- one of colours: green (0), **red (even number from 1 to 36)** and black (odd number from 1 to 36).

All wining bets are paid the odds plus the original stake. Some typical odds are:

i. 35 to 1 for a single number 0 to 36.
ii. 1 to 1 for red or black.
iii. 2 to 1 for dozens (1 to 12).

Roulette **Roulette-table**

What is the expected gain when you bet 1 on **red**?

5.1 Expectation

Example 2. A company of leather goods must decide whether to expand the capacity. The advisers provided the following information: if the company expands now and the **recession remains**, there will be a loss of **40,000** during the next fiscal year; but if the economic conditions will be improved, there will be a profit of 164,000. If we know that the probability of 'remain recession' is 2/3, what is the expected value of the profit in the next fiscal year?

Solutions. Let X be the profit in the next fiscal year, then

X	164,000	**-40,000**
p_i	1/3	**2/3**

Thus, $E(X) = 164,000 \times \frac{1}{3} + (-40,000) \times \frac{2}{3} = 28,000.$

5.1 Expectation

Example 3. If X has a uniform distribution $U(a, b)$, then it has the following density function

$$f(x) = \frac{1}{b-a}, \quad \text{for} \ \ a < x < b.$$

The expectation of X is

$$\begin{aligned} E(X) &= \int_{-\infty}^{\infty} xf(x)dx \\ &= \int_a^b \frac{x}{b-a}dx \\ &= \frac{1}{b-a}\left(\frac{b^2}{2} - \frac{a^2}{2}\right) = \frac{a+b}{2}. \end{aligned}$$

5.2 The decision model

Our concern is **decision**.

- Choose an action from an action space: $A = \{a_1, \ldots, a_n\}$. In Example 2, the action could be

$$a_1 = \{\text{expand now}\} \text{ or } a_1 = \{\text{delay expansion}\}.$$

- The consequence of the decision depends on the uncertainty regarding with the **true state of nature**.
  - The true state θ may be one of the outcome in the sample space $\Theta = (\theta_1, \theta_2, \ldots, \theta_m)$.
  - In Example 2, θ could be either θ_1 or θ_2:

$$\begin{aligned} \theta_1 &= \{\textbf{the economic conditions will be improved}\} \text{ and} \\ \theta_2 &= \{\textbf{remain recession}\}. \end{aligned}$$

- Once the decision has been made, the decision maker (engineer) can only wait to see which state of nature θ is the true one.

5.2 The decision model

Example 2 (continued) Denote actions as $a_1 = \{$**expand now**$\}$ and $a_2 = \{$delay expansion$\}$, and the state of nature as

$$\begin{aligned} \theta_1 &= \{\text{the economic conditions will be improved}\} \text{ and} \\ \theta_2 &= \{\text{remain recession}\}. \end{aligned}$$

- If we take action a_1, we have

$$u(a_1, \theta_1) = 164,000, \quad u(a_1, \theta_2) = -40,000.$$

- We further know that if the company waits at least another year to expand and the recession remains, there will be a profit of 8,000; and there will be a profit of 80,000 if the economic conditions will be improved. Thus, If we take action a_2,

$$u(a_2, \theta_1) = 80,000, \quad u(a_2, \theta_2) = 8,000.$$

Decision tree

5.2 The decision model

- A decision tree can be constructed for discrete-state space;
- **Utility function** $u(a, \theta)$: a numerical measure (say pounds) of the consequences of this action-**state** pair.
- Alternatively, we can define a **loss function** $l(a, \theta)$

5.2 The decision model

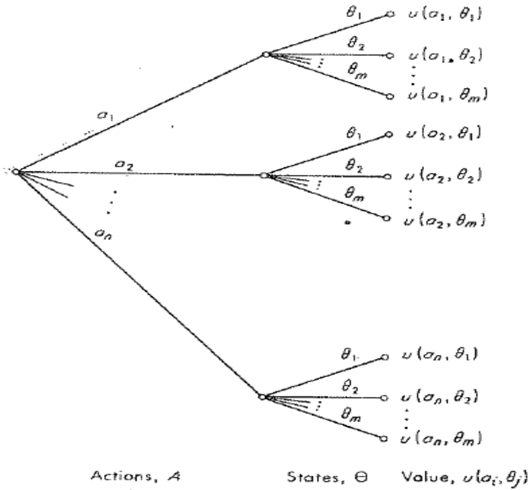

5.2 The decision model

Example 4. A construction engineer needs to select a steel pile length when the depth to rock is uncertain. The available actions are driving a 40- or 50-ft pile, and the possible states of nature are a **40- or 50-ft depth to bedrock**. The consequences of any action-state pair can be given in a **payoff** table:

	a_1 (drive 40-ft pile)	a_2 (drive 50-ft pile)
θ_1 depth 40 ft	correct decision no loss $l(a_1, \theta_1) = 0$	cut off 10-ft piece of pile 100-unit loss $l(a_2, \theta_1) = 100$
θ_2 depth 50 ft	splice and weld 10-ft pile 400-unit loss $l(a_1, \theta_2) = 400$	correct decision no loss $l(a_2, \theta_2) = 0$

Decision tree

5.2 The decision model: Example 4

In this example, we can define either a **loss function** $l(a_i, \theta_j)$ or a **utility function** $u(a_i, \theta_j)$:

loss function	utility function
$l(a_1, \theta_1) = 0$	$u(a_1, \theta_1) = 0$
$l(a_1, \theta_2) = 400$	$u(a_1, \theta_2) = -400$
$l(a_2, \theta_1) = 100$	$u(a_2, \theta_1) = -100$
$l(a_2, \theta_2) = 0$	$u(a_2, \theta_2) = 0$

5.2 The decision model

Example 5. Same as Example 4, the engineer can take action to drive either a 40- or 50-ft pile, but the possible states of nature θ has a uniform distribution $U(40, 50)$. We may consider the following loss function:

$$l(a_1, \theta) = 40 * (\theta - 40), \quad \theta \sim U(40, 50);$$

and

$$l(a_2, \theta) = 10 * (50 - \theta), \quad \theta \sim U(40, 50).$$

How to draw a decision tree?

Decision tree

5.3. Decision Theory: Minimax criterion

- For each action, we calculate the loss function (or utility function) in the worst-case scenario;
- We take the **action** to minimize the loss function (or maximize utility function).

5.3. Decision Theory: Minimax criterion

Example 2 (continued) If we don't know the probability of $\theta_1 = \{$**the economic conditions will be improved**$\}$ and $\theta_2 = \{$**remain recession**$\}$, what decision should be made?

- We can define the utility function as the profit achieved in the **worst-case scenario**. Thus

  $$\text{Utility}(\text{Expand now}) = u(a_1) = \min\{u(a_1, \theta_1), u(a_1, \theta_2)\} = -40,000,$$

  and

  $$\text{Utility}(\text{Delay expansion}) = u(a_2) = \min\{u(a_2, \theta_1), u(a_2, \theta_2)\} = 8,000.$$

- We should delay expansion by **maximising** the utility function for the worst case.
  Decision tree

5.3. Decision Theory: Minimax criterion

Example 4 (continued). It is quite obvious that we should take action a_2 based on Minimax criterion. Why?

Decision tree

5.3. Decision Theory: Expected-value decisions

How to make decision when a decision situation involves uncertainty?

- Let $u(a_i, \theta)$ be a utility function for action a_i — depends on a random variable θ (state of nature).
- Calculate the expected-value of utility function (or loss function).

5.3. Decision Theory: Expected-value decisions

- $u(a_i, \theta)$ — a utility function for action a_i.
- If θ is a discrete random variable with probability function $P(\theta = \theta_j) = p_j, j = 1, \ldots, m$, then the expected utility is

$$u(a_i) = E(u(a_i, \theta)) = \sum_{j=1}^{m} u(a_i, \theta_j) p_j;$$

- If θ is a continues random variable with density function $f(\theta)$, then the expected utility is

$$u(a_i) = E(u(a_i, \theta)) = \int u(a_i, \theta) f(\theta) d\theta.$$

- **Optimal decision** — choose the action with the maximum utility.

5.3. Decision Theory: Expected-value decisions

Example 2 (continued) From the discuss in Section 5.1, we have

$$
\begin{aligned}
P(\theta_1) &= P(\textbf{the economic conditions will be improved}) = 1/3 \\
P(\theta_2) &= P(\textbf{remain recession}) = 2/3
\end{aligned}
$$

- The expected value of the profit for action: **expand now**:

$$u(a_1) = E(u(a_1, \theta)) = 164,000 \times \frac{1}{3} + (-40,000) \times \frac{2}{3} = 28,000.$$

- The expected value for action: **delay expansion**

$$u(a_2) = E(u(a_2, \theta)) = 80,000 \times \frac{1}{3} + 8,000 \times \frac{2}{3} = 32,000.$$

- **Optimal decision** — **delay expansion**. **Decision tree**

5.3. Decision Theory: Expected-value decisions

Example 5 (continued). Since $\theta \sim U(40, 50)$, we know that (from Example 3)

$$E(\theta) = (40 + 50)/2 = 45.$$

- The expected values of the loss functions are:

$$\begin{aligned} l(a_1) &= E(l(a_1, \theta)) = 40 * (E(\theta) - 40) = 40 * (45 - 40) = 200. \\ l(a_2) &= E(l(a_2, \theta)) = 10 * (50 - E(\theta)) = 10 * (50 - 45) = 50. \end{aligned}$$

- **Optimal decision** — take action a_2 i.e., driving a 50-ft pile.
- The decision is made based on minimising the expected-value of loss function.

Example

A baker must decide to bake 100 or 150 loaves of bread on a given day. The cost of baking one loaf is 50 pence and the loaf sells for 1 pound. Assume that the demand on that day is 100 loaves with probability $2/3$ or 150 loaves with probability $1/3$.

1. Draw a decision tree.
2. Which plan would the baker use based on the minimax criterion?
3. Which plan would the baker use based on expected-value decisions?

5.3. Decision Theory: Bayesian decision

- It might be not easy to obtaine the probability function $P(\theta = \theta_j)$ or density function $f(\theta)$.
- Using Bayesian approach.
- If we have observed a set of data, we can calculate the posterior distribution

$$f(\theta|\text{data}).$$

- Calculate the expected-values of the utilities using the posterior distribution.