Lecture 8: Game theory

Dr. Roman V Belavkin

BIS3226

Contents

Historical Background

1654 Blaise Pascal and Pierre Fermat

1657 Christian Huygens publishes *On Ratiocination in Dice Games*

1944 John von Neumann and Oscar Morgenstern

1950 John Nash

1 Expected Utility and Decision-Making under Uncertainty

Reminder: Choice and Utility

- Choice set: $\Omega = \{a, b, c\}$, and preference relation $\lesssim$ on Ω:

$$a \lesssim b \lesssim c$$

- Utility function $u : \Omega \to \mathbb{R}$, $u(\omega_1) \leq u(\omega_2) \iff \omega_1 \lesssim \omega_2$:

$$-\pounds 10 \leq \pounds 0 \leq \pounds 10$$

- The optimal choice $c \in \Omega$ corresponds to the maximum of utility:

$$\max u(\omega) = u(c) = \pounds 10$$

Question 1. *How to make choice (or decisions) under uncertainty? That is if the elements of Ω appear with probability $P(\omega)$.*

Expected Utility

- Let $P(\omega)$ be a probability distribution over the elements of Ω:

$$P(\omega) = \{P(\omega_1), \dots, P(\omega_n)\}\,, \quad \sum_{i=1}^{n} P(\omega_i) = 1$$

- Thus, utility of $\omega \in \Omega$ is a random variable, and its expected values is the *expected utility*:

$$E_P\{u\} = u(\omega_1)P(\omega_1) + \cdots + u(\omega_n)P(\omega_n)$$

Example 1. $\Omega = \{a, b, c\}$, $P(\omega) \in \{0.1, 0.1, 0.8\}$, $u(\omega) \in \{-£10, £0, £10\}$

$$E_P\{u\} = -£10 \cdot 0.1 + £0 \cdot 0.1 + £10 \cdot 0.8 = £7$$

The Maximum Expected Utility Principle

- Under uncertainty, the choice between elements of Ω is replaced by a choice between different probability distributions on Ω:

$$P(\omega) = \{P(\omega_1), \dots, P(\omega_n)\}\,, \quad Q(\omega) = \{Q(\omega_1), \dots, Q(\omega_n)\}$$

- Probability distribution $P(\omega)$ is preferred to $Q(\omega)$ if and only if

$$E_P\{u\} \geq E_Q\{u\}$$

- In other words, the optimal choice is a distribution that yields the *maximum expected utility*:

$$\max_{P,\dots,Q} E_P\{u\}$$

Example: Maximum EU

Example 2. A lottery is a probability distribution over the prizes.

- Lottery A, in which you can win £1000, but also you may loose £100;
- Lottery B, in which you can win £100, but you can loose £10.
- Suppose also that the probability of winning in both lotteries is $\frac{1}{2}$

$$\begin{aligned} E_A\{u\} &= £1000\frac{1}{2} - £100\frac{1}{2} = £450 \\ E_B\{u\} &= £100\frac{1}{2} - £10\frac{1}{2} = £45 \end{aligned}$$

Thus, we prefer A to B.

2 Games

Games

- Games are conflict situations between several players (agents)

$$A,\ B,\ C,\ \ldots$$

- Each player has a set of strategies:

$$S^A = \{s_1^A, \ldots, s_m^A\}, \quad S^B = \{s_1^B, \ldots, s_n^B\}$$

- Each state of the game is defined by the strategies the players choose:

$$\Omega = S^A \times S^B \times \cdots$$

- Each player receives a utility (pay off) in each state:

$$u_A = u_A(s_i^A, s_j^B, \ldots)\,, \quad u_B = u_B(s_i^A, s_j^B, \ldots)$$

- Games can be finite, infinite, cooperative, non-cooperative, symmetric, asymmetric, zero-sum, non-zero-sum, sequential, simultaneous, etc.

Classical Game Examples

Example 3 (Rock Paper Scissors). For each player: $S = \{r, p, s\}$ with $r \lesssim p \lesssim s \lesssim r$ and $r \sim r$, $p \sim p$, $s \sim s$.

Example 4 (Penny Matching). Each player chooses $S = \{h, t\}$. Match if hh or tt; mismatch otherwise.

$$u_A(\text{mismatch}) \leq u_A(\text{match})\,, \quad u_B = -u_A$$

Example 5 (Prisoners' Dilemma). Two prisoners (players), each can $S = \{\text{cooperate}, \text{defect}\}$. Utility:

$$\begin{aligned}
u_A(\text{cooperate}, \text{cooperate}) &= 5\text{years} = u_B(\text{cooperate}, \text{cooperate}) \\
u_A(\text{defect}, \text{defect}) &= 0.5\text{years} = u_B(\text{defect}, \text{defect}) \\
u_A(\text{cooperate}, \text{defect}) &= 0 \\
u_A(\text{defect}, \text{cooperate}) &= 10\text{years}
\end{aligned}$$

3 Finite Zero-Sum 2-Person Games

Finite Zero-Sum 2-Person Games

- Two players A, B, each has finite set of strategies:

$$S^A = \{s_1^A, \ldots, s_m^A\}\,, \quad S^B = \{s_1^B, \ldots, s_n^B\}$$

- In a zero-sum game, the utility can be defined by the $m \times n$ *payoff matrix* of player A:

$$u_A = \begin{pmatrix} u_{11} & u_{12} & \cdots & u_{1n} \\ u_{21} & u_{22} & \cdots & u_{2n} \\ \cdot & \cdot & \cdots & \cdot \\ u_{m1} & u_{m2} & \cdots & u_{mn} \end{pmatrix}$$

where $u_{ij} = u_A(s_i^A, s_j^B)$.

- The payoff matrix of player B is $u_B = -u_A$, so that

$$u_A + u_B = 0$$

- If $m = n$, then the game is *symmetric*.

Finite Zero-Sum 2-Person Games (Examples)

Example 6 (Rock Paper Scissors).

$$u_A = \begin{pmatrix} u_{rr} & u_{rp} & u_{rs} \\ u_{pr} & u_{pp} & u_{ps} \\ u_{sr} & u_{sp} & u_{ss} \end{pmatrix} = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}, \quad u_B = -u_A$$

Example 7 (Penny Matching).

$$u_A = \begin{pmatrix} u_{hh} & u_{ht} \\ u_{th} & u_{tt} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \quad u_B = -u_A$$

Example 8 (Prisoners' Dilemma).

$$u_A = \begin{pmatrix} u_{cc} & u_{cd} \\ u_{dc} & u_{dd} \end{pmatrix} = \begin{pmatrix} -5 & 0 \\ -10 & -.5 \end{pmatrix}, \quad u_B \neq -u_A$$

Thus, prisoners' dilemma is *not* a zero-sum game.

Max-Min and Min-Max Principle in Zero-Sum Games

- Each player can only choose their own strategies.
- Player A chooses s^A to maximise $\min_{s^B} u_A$.
- Player B chooses s^B to maximise $\min_{s^A} u_B$.
- Because $u_B = -u_A$, the latter is equivalent to player B chooses s^B to minimise $\max_{s^A} u_A$.
- For any payoff matrix $u_A = (u_{ij})$, the following is true

$$\max_i \min_j u_{ij} \leq \min_j \max_i u_{ij}$$

- The game has a *solution* (or a *saddle point*), if there exists a pair ij (i.e. a pair of strategies s_i^A and s_j^B) such that

$$\max_i \min_j u_{ij} = \min_j \max_i u_{ij}$$

4 Mixed Strategies

Mixed Strategies

Definition 9 (Mixed Stragtegy)**.** of player A is a probability distribution $P^A = \{p_1^A, \dots, p_n^A\}$ over the set of his strategies $S^A = \{s_1^A, \dots, s_m^A\}$.

Example 10 (Rock Paper Scissors)*.* $P^A = \{p_r^A, p_p^A, p_s^A\}$. A mixed strategy such that $p_i^A = 1/3$ for all $i \in \{r, p, s\}$ chooses rock, paper or scissors with equal probability.

Expected payoff
If P^A and P^B are mixed strategies of players A and B, then the expected utility (payoff) to players A and B are:

$$E_{P^A P^B}\{u_A\} = \sum_{i=1}^{m} \sum_{j=1}^{n} u_{ij} p_i^A p_j^B \,, \qquad E_{P^A P^B}\{u_B\} = -E_{P^A P^B}\{u_A\}$$

Mixed Strategies (Examples)

Example 11 (Penny Matching)*.*

- Strategies $S^A = S^B = \{h, t\}$. Let mixed strategies be:

$$P^A = \{.8, .2\}\,, \quad P^B = \{.6, .4\}$$

- Given payoff matrix

$$u_A = \begin{pmatrix} u_{hh} & u_{ht} \\ u_{th} & u_{tt} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

- The expected payoff is

$$\begin{aligned} E_{P^A P^B}\{u_A\} &= 1 \cdot 0.8 \cdot 0.6 - 1 \cdot 0.2 \cdot 0.6 - 1 \cdot 0.8 \cdot 0.4 + 1 \cdot 0.2 \cdot 0.4 \\ &= 0.48 - 0.12 - 0.32 + 0.08 = 0.12 \\ E_{P^A P^B}\{u_B\} &= -0.12 \end{aligned}$$

Max-Min and Min-Max Solutions for Mixed Strategies

- Each player can only choose their own mixed strategies P^A and P^B.
- Player A chooses P^A to maximise $\min_{P^B} E_{P^A P^B}\{u_A\}$.
- Player B chooses P^B to minimise $\max_{P^A} E_{P^A P^B}\{u_A\}$.
- Every finite zero-sum 2-person game has a *solution* (or a *saddle point*), defined by the optimal mixed strategy $\bar{P}^A \bar{P}^B$ such that

$$\max_{P^A} \min_{P^B} E_{P^A P^B}\{u_A\} = \min_{P^B} \max_{P^A} E_{P^A P^B}\{u_A\}$$

- The common value $E_{\bar{P}^A \bar{P}^B}\{u_A\}$ is called the *value* of the game.
- A zero-sum 2-person game is *fair* if its value is zero.

Summary

- Games are models of conflict situations between several players.
- Preference relations on the states of the game and the corresponding utility functions define the payoff matrix.
- The expected utility allows us to find optimal mixed (probabilistic) strategies.