IB Optimisation: Lecture 9

Mike Tehranchi

University of Cambridge

13 May 2020

Consider a game

- there are two players, called Player I and Player II
- Player I has m choices of strategies, labelled $i \in \{1, \dots, m\}$
- Player II has n choices of strategies, labelled $j \in \{1, \dots, n\}$.
- The key assumption is that the game is zero-sum: If Player I chooses strategy i and Player II chooses strategy j, then
  - Player I is paid £$a_{i,j}$
  - Player II is paid £$(-a_{i,j})$.

  In particular, the net payment is zero.
- The matrix $A = (a_{i,j})_{i,j}$ is called the *payoff* matrix of the game.

Preliminary analysis.

- Player I wants to maximise his payoff
- but he knows that Player II wants to minimise it
- Player I might want to solve

$$\text{maximise } \min_j a_{ij} \text{ subject to } i \in \{1, \ldots, m\}.$$

- Similarly, Player II might want to solve

$$\text{minimise } \max_i a_{ij} \text{ subject to } j \in \{1, \ldots, n\}.$$

Example. Consider a game with payoff matrix

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$$

$i \backslash j$			row min	
	1	2	1	
	$\boxed{3}$	4	3	$\leftarrow$
col max	3	4		
	$\uparrow$			

- Player I will pick $i = 2$
- Player II will pick $j = 1$.
- The point $(2, 1)$ is called a saddle point of the matrix.

Definition

A *saddle point* of a payoff matrix A is a pair of strategies (i, j) such that

$$a_{i,j} = \max_{i'} \min_{j'} a_{i',j'} = \min_{j'} \max_{i'} a_{i',j'}.$$

If a payoff matrix A has a saddle point (i, j), then the element $a_{i,j}$ is called the *value of the game*.

Not all payoff matrices have a saddle point. Here is an example:

$i \backslash j$			row min	
	4	2	2	$\leftarrow$
	1	3	1	
col max	4	3		
		$\uparrow$		

- If Player I was to play first, he would play row 1 since this would maximise the minimum of $a_{i,j}$, yielding payoff $\max_i \min_j a_{i,j} = 2$.
- If Player II was to play first, she would pick column 2, since this would minimise the maximum of $a_{i,j}$. For Player II, the loss woud be $\min_j \max_i a_{i,j} = 3$.

What if the two players pick their strategies simultaneously? Idea: randomised strategies.

Definition

A *mixed strategy* is an assignment of probabilities to each of the individual strategies. A *pure* strategy is a mixed strategy that assigns probability 1 to one of the strategies and 0 to the rest.

From now on we allow the players to use mixed strategies, and use the notation

- Player I plays strategy i with probability p_i, and
- Player II plays strategy j with probability q_j.

Player I wants to

$$\text{maximise } \min_j \mathbb{E}(\text{ payout} \mid \text{Player II picks strategy } j).$$

In notation, this is to

$$\text{maximise } \min_j \sum_{i=1}^{m} p_i a_{i,j} \text{ subject to } \sum_{i=1}^{m} p_i = 1 \text{ and } p_i \geq 0 \text{ for all } i$$

We can turn this into a linear program. Using the notation $e = (1, \ldots, 1)^\top$ for a column vector in $\mathbb{R}^m$ or $\mathbb{R}^n$ as context dictates, Player I's problem is

$$\text{I: maximise } v \text{ subject to } A^\top p \geq ve, \; e^\top p = 1, \; p \geq 0.$$

Similarly, Player II's problem is

$$\text{II: minimise } \max_i \sum_{j=1}^{n} a_{i,j} q_j \text{ subject to } \sum_{j=1}^{n} q_j = 1 \text{ and } q_j \geq 0 \text{ for all } j$$

or equivalently, to the linear program to

$$\text{II: minimise } w \text{ subject to } Aq \leq we, \; e^\top q = 1, \; q \geq 0.$$

Player II's linear program is the dual of Player I's.
By the fundamental theorem of linear programs that Player I's solution (p, v) is optimal if and only if there exists (q, w) such that

- (p, v) is feasible for the primal problem,
- (q, w) is feasible for the dual problem, and
- the solutions are matched by complementary slackness

$$(Aq - we)^\top p = 0 = q^\top (A^\top p - ve).$$

Theorem (Fundamental theorem of matrix games)

The mixed row strategy p is optimal for Player I's problem if and only if there exists a mixed column strategy q and $v \in \mathbb{R}$ such that

- $A^\top p \geq ve, \; e^\top p = 1, \; p \geq 0.$ *(primal feasibility)*
- $Aq \leq we, \; e^\top q = 1, \; q \geq 0$ *(dual feasibility)*
- $v = p^\top Aq$ *(complementary slackness).*

In this case, (p, v) is optimal for Player I's linear program, (q, v) is optimal for Player II's linear program, and the quantity v is called the value of the game.

The complementary slackness condition means that for optimal mixed strategies:
If II plays j with positive probability, then the conditional expected payoff given II plays j equals the value of the game.
Or in notation

$$q_j > 0 \Rightarrow (A^\top p)_j = v.$$

Similarly, we have

$$p_i > 0 \Rightarrow (Aq)_i = v.$$

Definition

A game is *symmetric* if $m = n$ and the payoff matrix $A = -A^\top$ is anti-symmetric.

Example: Rock-paper-scissors

	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

Theorem

The value of a symmetric game is zero.

Proof. Suppose the triple (p, v, q) satisfy the (necessary) conditions of optimality. Since $A = -A^{\top}$, the triple $(q, -v, p)$ satisfies the (sufficient) conditions of optimality. Thus $v = -v$. □

(1) *Look for a saddle point.* If a saddle point (i, j) exists, then the optimal strategy for player I is the pure strategy i, and the optimal strategy for player II is the pure strategy j. The value of the game is $a_{i,j}$.

(2) *Look for dominating strategies.*

- Row i *dominates* row i' if

$$a_{i,j} \geq a_{i',j} \text{ for all } j = 1, \ldots, n$$

  Player I should never play strategy i'.

- Similarly, column j *dominates* column j' if

$$a_{i,j} \leq a_{i,j'} \text{ for all } i = 1, \ldots, m.$$

  Player II never plays strategy j'.

Example. Consider

$$A = \begin{pmatrix} 2 & 3 & 4 & 2 \\ 3 & 1 & 1/2 & 4 \\ 1 & 3 & 2 & 3 \end{pmatrix}$$

- There is no saddle point.
- Column 1 dominates column 4.
- After eliminating column 4, row 1 dominates row 3
- Player I's optimal strategy p is of the form $= (p_1, 1 - p_1, 0)^\top$.

(3) *Draw a picture.* Consider

$$A' = \begin{pmatrix} 2 & 3 & 4 \\ 3 & 1 & 1/2 \end{pmatrix}.$$

The region $(A')^\top p' \geq ve$ where $p' = (p_1, 1 - p_1)^\top$ is given by

$$\begin{aligned} 2p_1 + 3(1 - p_1) \geq v &\Rightarrow v \leq 3 - p_1 \\ 3p_1 + (1 - p_1) \geq v &\Rightarrow v \leq 1 + 2p_1 \\ 4p_1 + \tfrac{1}{2}(1 - p_1) \geq v &\Rightarrow v \leq \tfrac{1}{2} + \tfrac{7}{2}p_1. \end{aligned}$$

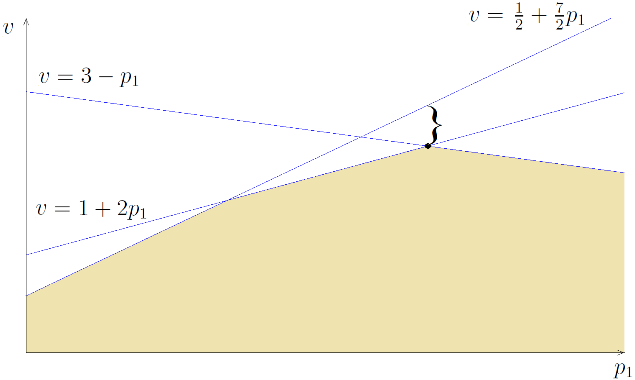

- maximum v occurs at $p_1 = 2/3$.
- Player I's optimal strategy is $p = (2/3, 1/3, 0)^\top$
- the value of the game is $v = 7/3$.

To find Player II's optimal strategy,

- At the point $(p_1, v) = (2/3, 7/3)$ the third constraint $v < \frac{1}{2} + \frac{7}{2}p_1$ is not binding.
- By complementary slackness $q_3 = 0$. So the optimal q is of the form $q = (q_1, 1 - q_1, 0, 0)^\top$.
- Since $p_1 = 2/3$ is positive, by complementary slackness, the first dual constraint is binding, yielding

$$2q_1 + 3(1 - q_2) = 7/3 \Rightarrow q_1 = 2/3$$

(4) *Use the simplex algorithm.* If $\min_{i,j} a_{i,j} > 0$ we know that the value of the game is strictly positive. Hence we can put the problem in a form to use the simplex algorithm as follows:

- Let $x = p/v$ so that Player I's problem becomes

$$\text{maximise } v \text{ subject to } A^\top x \geq e, e^\top x = 1/v, \; x \geq 0$$

- equivalent to

$$\text{minimise } e^\top x \text{ subject to } A^\top x \geq e, \; x \geq 0$$

- We could use the two-phase method.
- Or we could look at the dual problem

$$\text{maximise } e^\top y \text{ subject to } Ay \leq e, y \geq 0$$

which is exactly in the form to use the one-phase method.

If $\min_{i,j} a_{i,j} \leq 0$ we can still use the above idea.

- Find a k such that $a'_{i,j} = a_{i,j} + k > 0$ for all i, j.
- Then solve the problem for this new payout matrix.
- The optimal strategies p, q of both the original and modified games will be the same
- The value v of the original game can be calculated from the value v' of the modified game by $v = v' - k$.

Example: Rock-paper-scissors again. Take $k = 2$.

$$A' = \begin{pmatrix} 2 & 1 & 3 \\ 3 & 2 & 1 \\ 1 & 3 & 2 \end{pmatrix}$$

Player II's problem is equivalent to

$$\text{maximise } y_1 + y_2 + y_3 \text{ subject to } \begin{array}{ll} 2y_1 + y_2 + 3y_3 & \leq 1 \\ 3y_1 + 2y_2 + y_3 & \leq 1 \\ y_1 + 3y_2 + 2y_3 & \leq 1 \\ y_1, y_2, y_3 & \geq 0 \end{array}$$

	*	*	*				
	y_1	y_2	y_3	z_1	z_2	z_3	
z_1	2	1	3	1	0	0	1
z_2	3	2	1	0	1	0	1
z_3	1	3	2	0	0	1	1
Payoff	1	1	1	0	0	0	0

	y_1	y_2	y_3	z_1	z_2	z_3	
	*	*	*				
y_3	0	0	1	$\frac{7}{18}$	$-\frac{5}{18}$	$\frac{1}{18}$	$\frac{1}{6}$
y_1	1	0	0	$\frac{1}{18}$	$\frac{7}{18}$	$-\frac{5}{18}$	$\frac{1}{6}$
y_2	0	1	0	$-\frac{5}{18}$	$\frac{1}{18}$	$\frac{7}{18}$	$\frac{1}{6}$
Payoff	0	0	0	$-\frac{1}{6}$	$-\frac{1}{6}$	$-\frac{1}{6}$	$-\frac{1}{2}$