Maximin Strategies in Zero-Sum Games

- If a player adopts a maximin strategy, he believes that the opponent will always correctly predict their move.
- This means, the opponent will choose their best possible action based upon the player's act.
- In this case, player 1's expected payoff is:

$$R_{\text{maximin}}(d_i) = \min_j R(d_i, \delta_j)$$

- If this is the case, then player 2's payoff is:

$$-R_{\text{maximin}}(d_i) = \max_j -R(d_i, \delta_j)$$

- Hence $P1$ should play $d^\star_{\text{maximin}} = \arg\max_{d_i} \min_j R(d_i, \delta_j)$.
- One could swap the two players to obtain a maximin strategy for player 2.

Example (RPS and Maximin)

- Let $M = (m_{ij})$ denote the payoff matrix for the RPS game.
- Then, $\min_j R(d_i, \delta_j) = \min_j m_{ij} = -1$ for all i.
- Thus any move is maximin for player 1.
- Player 1 expects to receive a payout of -1 whatever he does.
- If both players adopt a maximin view, then player 2 has the same expectation (by symmetry).
- How can we resolve this paradox?

What's Gone Wrong?

- The players aren't using all of the information available.
- They haven't used the fact that it is a zero sum game.
- They don't have compatible beliefs:
  - If P1 believes P2 can predict their move and P2 believes that P1 can predict their move then things inevitably go wrong.
  - It cannot be common knowledge that *both* players will adopt a maximin strategy!
- If a player really believes their opponent can predict their move then they can use randomization to make their action less predictable. . .

Mixed Strategies

- A *mixed strategy* for player 1 is a probability distribution over D.
- If a player has mixed strategy $\mathbf{x} = (x_1, \ldots, x_n)$ then they will play move d_i with probability x_i.
- This can be achieved using a randomization device such as a spinner to select a move.
- A *pure* strategy is a mixed strategy in which exactly one of the x_i is non-zero (and is therefore equal to 1).
- A similar definition applies when considering player 2.

Expected Rewards and Mixed Strategies

What is player 1's expected reward if. . .

- Player 1 has mixed strategy $\underline{x}$ and player 2 plays pure strategy δ_j?
- Player 1 has pure strategy d_i and player 2 plays mixed strategy $\underline{y}$?
- Player 1 has mixed strategy $\underline{x}$ and player 2 has mixed strategy $\underline{y}$?

In the first case, the uncertainty is player 1's own move, and his expectation is:

$$\sum_{i=1}^{n} x_i R(d_i, \delta_j)$$

In the second case, the uncertainty comes from player 2:

$$\sum_{j=1}^{m} y_j R(d_i, \delta_j)$$

Whilst both provide (independent) uncertainty in the third case:

$$\sum_{i=1}^{n} \sum_{j=1}^{m} x_i R(d_i, \delta_j) y_j = \underline{x}^{\mathsf{T}} M \underline{y}$$

Maximin Revisited

- Player 1's maximin *mixed* strategy is the $\underline{x}$ which maximises:

$$V_1 = \max_{\underline{x}} \min_{\underline{y}} \sum_i \sum_j x_i R(d_i, \delta_j) y_j$$

- Player 2's maximin *mixed* strategy is the $\underline{y}$ which minimises:

$$\max_{\underline{y}} \min_{\underline{x}} -\sum_i \sum_j x_i R(d_i, \delta_j) y_j$$

$$= \min_{\underline{y}} \max_{\underline{x}} \sum_i \sum_j x_i R(d_i, \delta_j) y_j = V_2$$

What is the relationship between these two values?

Theorem (Fundamental Theorem of Zero Sum Two Player Games)

V_1 and V_2 as defined before satisfy:

$$V_1 = V_2$$

The unique value, $V = V_1 = V_2$ is known as the value of the game.

- The strategies $\underline{x}$ and $\underline{y}$ which achieve this value may not be unique.
- How can we find suitable strategies in general?

(Sketch of proof of theorem see later.)

Example (Maximin in a Simple Game)

- Consider a zero sum two player game with the following payoff matrix:

	δ_1	δ_2
d_1	1	3
d_2	4	2

- With a pure strategy maximin approach:
  - P1 plays d_2 expecting P2 to play δ_2.
  - P2 plays δ_2 expecting P1 to play d_1.
  - P1 expects to gain 2; P2 expects to lose 3.
  - This is not consistent.

Example

- Consider, instead, a mixed strategy maximin approach:
  - P1 plays a strategy $(x, 1-x)$ and player 2 plays $(y, 1-y)$.
  - Player 1's expected payoff is:

(= - 4xy + x + 2y + 2, use algebra…)

$$[x \quad 1-x] \begin{bmatrix} 1 & 3 \\ 4 & 2 \end{bmatrix} \begin{bmatrix} y \\ 1-y \end{bmatrix} = -4(x - \frac{1}{2})(y - \frac{1}{4}) + \frac{5}{2}$$

  - Player 1 seeks to maximise this for the worst possible y.
  - As the 2nd player can control the sign of the first term, his optimal strategy is to make it vanish by choosing $x = \frac{1}{2}$.
  - Similarly, the 2nd player wants to prevent the first player from exploiting the first term and chooses $y = \frac{1}{4}$.
  - Now, the expected reward for the first player is, consistently, 2.5 as both expect the same maximin strategies to be played.
  - *Both* players have a higher expected return than they would playing pure strategies.

How do we determine maximin mixed strategies?

- We need a general strategy for determining strategies $\underline{x}^\star$ and $\underline{y}^\star$ which achieve the common maximin return for player 1.
- It's straightforward (if possibly tedious) to calculate, for payoff matrix M the expected return for player 1 as a function of the strategies:

$$V(\underline{x}, \underline{y}) = \underline{x}^\top M \underline{y}$$

- We then seek to obtain $\underline{x}^\star, \underline{y}^\star$ such that:

$$V(\underline{x}^\star, \underline{y}^\star) = \max_{\underline{x}} \min_{\underline{y}} V(\underline{x}, \underline{y})$$

- In general, this is a problem which can be efficiently addressed by linear programming.
- If one player has only two possible decisions, however, a simple graphical method can be employed. *(Only 1 parameter!)*

Graphical Solution, Part 1: Player 1's approach

- Consider a two player zero sum game with payoff matrix:

$$M = \begin{bmatrix} 2 & 3 & 11 \\ 7 & 5 & 2 \end{bmatrix}$$

- Consider a mixed strategy $(x, 1-x)$ for player 1.
- For the three pure strategies available to player 2, player 1 has expected reward:
  - $\delta_1 : 2x + 7(1-x) = 7 - 5x$
  - $\delta_2 : 3x + 5(1-x) = 5 - 2x$
  - $\delta_3 : 11x + 2(1-x) = 2 + 9x$
- For each value of x, the worst case response of player 2 is the one for which the expected reward of player 1 is minimised.
- Plotting the three lines as a function of x…

Zero-Sum Games

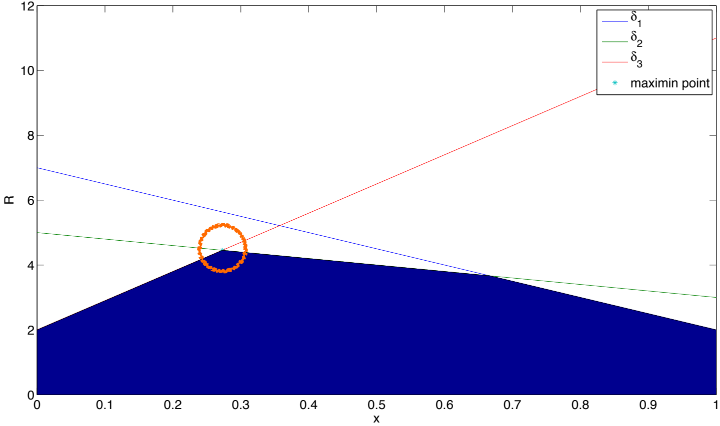

- The maximin response maximises the return in the worst case.
- In terms of our graph, this means we choose x to maximise the distance between the lowest of the lines and the ordinate axis.
- This is at the point where the lines associated with δ_2 and δ_3 intersect, at $x^\star$ which solves:

$$\begin{aligned} 5 - 2x =& 2 + 9x \\ 11x =& 3 \Rightarrow x^\star = 3/11 \end{aligned}$$

- Hence player 1's maximin mixed strategy is $(3/11, 8/11)$.
- Playing this, his expected return is:

$$V_1 = 2 + 9 \times 3/11 = 49/11 = \quad 5 - 2 \times 3/11 = 49/11$$

Graphical Solution, Part 2: Player 2's approach

- Player 2 only needs to consider the moves which optimally oppose player 1's maximin strategy (δ_2 and δ_3).
- They may consider a mixed strategy $(0, y, 1-y)$.
- By the fundamental theorem, player 2's maximn strategy leads to the same expected payoff for player 1 as his own maximin strategy:

$$V_2 = V_1 = 49/11.$$

2	3	11
7	5	2

- They should play $y^\star$ to solve:

$$V_2 = 3y + 11(1-y) = 49/11$$

$$8y = (121-49)/11 = 72/11 \Rightarrow y^\star = 9/11$$

- Leading to a mixed strategy $(0, 9/11, 2/11)$.

Example (Spy Game)

- A spy has escaped and must choose to flee down a *river* or through a *forest*. Their guard must choose to chasse them using a *helicopter*, a pack of *dogs* or a *jeep*.
- They agree that the probabilties of escape are as given in this payoff matrix:

	H	D	J
R	0.1	0.8	0.4
F	0.9	0.1	0.6

- Both players wish to adopt maximin strategies.

Example

- The spy plays strategy $(x, 1-x)$: with probability x they escape via the river; with probability $1-x$ they run through the forest.
- For given x, their probabilities of escaping for each of the guard's possible actions are:

$$p_H = 0.1x + 0.9(1-x) = \frac{9-8x}{10} \qquad p_D = 0.8x + 0.1(1-x) = \frac{1+7x}{10}$$

$$p_J = 0.4x + 0.6(1-x) = \frac{6-2x}{10}$$

- Plotting these three lines as a function of x we obtain the following figure:

Zero-Sum Games

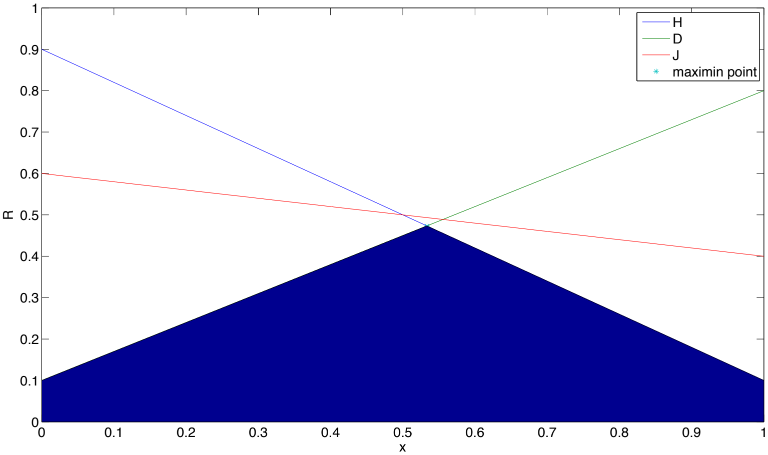

Example

- The maximin solution is the interesection of the lines for strategies D and H.
- This occurs at the solution, $x^\star$ of:

$$
\begin{aligned}
p_H = p_D \Rightarrow 9 - 8x =& 1 + 7x \\
8 =& 15x \qquad \Rightarrow x^\star = 8/15
\end{aligned}
$$

- The value of the game is: $V = V_1 = \frac{9-8x^\star}{10} = 71/150$

Example

- By the fundamental theorem of zero sum two player games, the guard needs to consider only H and D.
- Otherwise the spy's chance of escape will be better than V_1 if he plays his own maximin strategy.
- Consider a strategy $(y, 1-y, 0)$.
- By the same theorem, $V_2 = V = V_1$, so:

$$\begin{aligned} V_2 = 0.1y^\star + 0.8(1-y^\star) &= 71/150 \\ 8 - 7y^\star &= 71/15 \\ y^\star &= 7/15 \end{aligned}$$

On Zero Sum Two Player Games

- The “fundamental theorem” does not generalise to games of more than two players.
- The “fundamental theorem” does not generalise to non-zero-sum games.
- Games with an element of co-operation are much more interesting.