Games

What is a Game

A *game* in mathematics is, roughly speaking, a problem in which:

- Several *agents* or *players* make 1 or more decisions.
- Each player has an objective / set of preferences.
- The outcome is influenced by the set of decisions.
- There may be additional non-deterministic uncertainty.
- The players may be in competition or they may be cooperating.
- Examples include: chess, poker, bridge, rock-paper-scissors and many others.

However, we will stick to simple two player games with each player simultaneously making a single decision.

Simple Two Player Games

- Player 1 chooses a move for a set $D = \{d_1, \ldots, d_n\}$.
- Plater 2 chooses a move from a set $\Delta = \{\delta_1, \ldots, \delta_m\}$.
- Each player has a *payoff function*.
- If the players choose moves d_i and δ_j, then:
  - Player 1 receives reward $R(d_i, \delta_j)$.
  - Player 2 receives reward $S(d_i, \delta_j)$.
- The relationship between decisions and rewards is often shown in a payoff matrix:

	δ_1	$\ldots$	δ_m
d_1	$(R(d_1, \delta_1), S(d_1, \delta_1))$	$\ldots$	$(R(d_1, \delta_m), S(d_1, \delta_m))$
$\vdots$			$\vdots$
d_n	$(R(d_n, \delta_1), S(d_n, \delta_1))$	$\ldots$	$(R(d_n, \delta_m), S(d_n, \delta_m))$

Payoff Matrices Again

It's sometimes useful to consider a single player's payoff as a function of the possible decisions.

Player 1 and player 2 have these payoff matrices:

	δ_1	$\dots$	δ_m
d_1	$R(d_1, \delta_1)$	$\dots$	$R(d_1, \delta_m)$
$\vdots$			$\vdots$
d_n	$R(d_n, \delta_1)$	$\dots$	$R(d_n, \delta_m)$

	δ_1	$\dots$	δ_m
d_1	$S(d_1, \delta_1)$	$\dots$	$S(d_1, \delta_m)$
$\vdots$			$\vdots$
d_n	$S(d_n, \delta_1)$	$\dots$	$S(d_n, \delta_m)$

Example (Rock-Paper-Scissors)

- Each player picks from the same set of decisions:

$$D = \Delta = \{R, P, S\}$$

- R beats S; S beats P and P beats R
- One possible payoff matrix is:

	R	P	S
R	(0,0)	(-1,1)	(1,-1)
P	(1,-1)	(0,0)	(-1,1)
S	(-1,1)	(1,-1)	(0,0)

Example (The Prisoner's Dilemma)

- Again, each player picks from the same set of decisions:

$$D = \Delta = \{\text{Stay Silent}, \text{Betray Partner}\}$$

- If they both stay silent they will receive a short sentence; if they both betray one another they will get a long sentence; if only one betrays the other the traitor will be released and the other will get a long sentence.
- One possible payoff matrix is:

	S	B
S	(1,1)	(5,0)
B	(0,5)	(4,4)

- Notice that each player wishes to minimise this payoff!

Example (Love Story)

- A boy and a girl must go to either of:

$$D = \Delta = \{\text{Football}, \text{Opera}\}$$

- They both wish to meet one another most of all.
- If they don't meet, the boy would rather see the football; the girl, the opera.
- A possible payoff matrix might be:

	F	O
F	(100,100)	(50,50)
O	(0,0)	(100,100)

Some Features of these Examples

- The rock-paper-scissors game is *purely competitive*: any gain by one player is matched by a loss by the other player.
- The RPS and PD problems are symmetric:

$$R(d, \delta) = S(\delta, d)$$

  [Note that this makes sense as $D = \Delta$]
- $D = \Delta$ in all three of these examples, but it isn't always the case.

Uncertainty in Games

As the players don't know what action the other will take, there is uncertainty.

- Thankfully, the Bayesian interpretation of probability allows them to encode their uncertainty in a probability distribution.
- Player 1 has a probability mass function p over the actions that player 2 can take, Δ.
- Player 2 has a probability mass function q over the actions that player 1 can take, denoted D.

Expected Rewards

Just as in a decision problem, we can think about expected rewards:

- For player 1, the expected reward of move d_i is:

$$\bar{R}(d_i) = \mathbb{E}\left[R(d_i, \delta_j)\right] = \sum_{j=1}^{m} q(\delta_j) R(d_i, \delta_j)$$

- Whilst, for player 2, we have

$$\bar{S}(\delta_j) = \mathbb{E}\left[S(d_i, \delta_j)\right] = \sum_{i=1}^{n} p(d_i) S(d_i, \delta_j)$$

Some Interesting Questions

- When can a player act without considering what the opponent will do? i.e. When is player 1's strategy independent of p or player 2's of q?
- When p or q is important, how can rationality of the opponent help us to elicit them?
- What are the implications of this?

Separable Games

If we can decompose the rewards appropriately, then there is no interaction between the players' decisions:

- A game is *separable* if:

$$
\begin{aligned}
R(d,\delta) =& r_1(d) + r_2(\delta) \\
S(d,\delta) =& s_1(d) + s_2(\delta)
\end{aligned}
$$

- Here, the effect of the other player's act on a player's reward doesn't depend on their own decision:

$$
\begin{aligned}
\bar{R}(d_i) =& r_1(d_i) + \sum_{j=1}^{m} q(\delta_j) r_2(\delta_j) \\
\bar{S}(\delta_j) =& \sum_{i=1}^{n} p(d_i) r_1(d_i) + r_2(\delta_j)
\end{aligned}
$$

Strategy in Separable Games

- Player 1's strategy should depend only upon r_1 as the decision they make doesn't alter the reward from r_2.
- Player 2's strategy should depend only upon s_2 as the decision they make doesn't alter the reward from s_1.
- So, player 1 should choose a strategy from the set:

$$D^\star = \{d^\star : r_1(d^\star) \geq r_1(d_i) \quad i = 1, \ldots, n\}$$

- And player 2 from:

$$\Delta^\star = \{\delta^\star : s_2(\delta^\star) \geq s_2(\delta_j) \quad j = 1, \ldots, m\}$$

The Prisoner's Dilemma is a Separable Game

- Let $r_1(S) = 0$ and $r_1(B) = 1$.
- Let $r_2(S) = -1$ and $r_2(B) = -5$.
- Now, $R(d, \delta) = r_1(d) + r_2(\delta)$.
- And $D^\star = \{B\}$.
- Similarly for the second player, $\Delta^\star = \{B\}$.
- This is the so-called paradox of the prisoner's dilemma: both players acting rationally and independently leads to the worst possible solution!

Rationality and Games

As in decision theory, a rational player should maximise their expected utility. We will generally assume that utility is equal to payoff; no greater complications arise if this is not the case.

- For a given pmf q, player 1 has:

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

- Whilst for given p, player 2 has:

$$\bar{S}(\delta_j) = \sum_{i=1}^{n} S(d_i, \delta_j) p(d_i)$$

- We want p and q to be consistent with the assumption that the opponent is rational.
- We assume, that *rationality of all players is common knowledge.*

Common Knowledge: A Psychological Infinite Regress

In the theory of games the phrase *common knowledge* has a very specific meaning.

- Common knowledge is known by all players.
- That common knowledge is known by all players is known by all players.
- That common knowledge is common to all players is known by all players

  ⋮
- More compactly: common knowledge is something that is known by all players and the fact that this thing is known by all players is itself common knowledge.
- This is an example of an infinite regress.

Domination

- A move $d^\star$ is said to dominate all other strategies if:

$$\forall d_i \neq d^\star, j : \qquad R(d^\star, \delta^j) \geq R(d_i, \delta_j)$$

- It is said to *strictly dominate* those strategies if:

$$\forall d_i \neq d^\star, j : \qquad R(d^\star, \delta^j) > R(d_i, \delta_j)$$

- A move d' is said to be *dominated* if:

$$\exists i \text{ such that } d_i \neq d' \text{ and } \forall j : R(d', \delta_j) \leq R(d_i, \delta_j)$$

- It is said to be *strictly dominated* if:

$$\exists i \text{ such that } d_i \neq d' \text{ and } \forall j : R(d', \delta_j) < R(d_i, \delta_j)$$

Theorem (Dominant Moves Should be Played)

If a game has a payoff matrix such that player 1 has a dominant strategy, $d^\star$ then the optimal move for player 1 is $d^\star$ irrespective of q.

Proof:

- *Player 1 is rational and hence seeks the d_i which maximises*

$$\sum_j R(d_i, \delta_j) q(d_j)$$

- *Domination tells us that* $\forall i, j : \quad R(d^\star, \delta_j) \geq R(d_i, \delta_j)$
- *And hence, that:*

$$\sum_j R(d^\star, \delta_j) q(d_j) \geq \sum_j R(d_i, \delta_j) q(d_j)$$

A similar results holds for player 2.

Rationality and Domination

If rationality is common knowledge and $d^\star$ is a strictly dominant strategy for player 1 then:

- Player 1, being rational, plays move $d^\star$.
- Player 2, knows that player 1 is rational, and hence knows that he will play move $d^\star$.
- Player 2 can exploit this knowledge to play the optimal move *given that player 1 will play* $d^\star$.
- Player 2 plays moves $\delta^\star$ with $\delta^\star$ such that:

$$\forall j : S(d^\star, \delta^\star) \geq S(d^\star, \delta_j)$$

- If there are several possible $\delta^\star$ then one may be chosen arbitrarily.

Example (A game with a dominant strategy)

Consider the following payoff matrix:

	δ_1	δ_2	δ_3	δ_4
d_1	(2,-2)	(1,-1)	(10,-10)	(11,-11)
d_2	(0,0)	(-1,1)	(1,-1)	(2,-2)
d_3	(-3,3)	(-5,5)	(-1,1)	(1,-1)

- If rational, player 1 must choose d_1.
- Player 2 knows that player 1 will choose d_1.
- Consequently, player 2 will choose δ_2.
- (d_1, δ_2) is known as a discriminating solution.

Iterated Strict Domination

1. Let $D_0 = D$ and $\Delta_0 = 0$. Let t = 1
2. Player 1 checks D_{t-1} to see if it contains one or more strictly dominated moves. Let D'_t be the set of such moves.
3. Let $D_t = D_{t-1} \setminus D'_t$.
4. Player 1 checks D_{t-1} to see if it contains one or more strictly dominated strategies given that player 2 must choose a move from Δ_{t-1}. Let D'_t be the set of these strategies. Let $D_t = D_{t-1} \setminus D'_t$.
5. Player 2 updates Δ_{t-1} in the same way noting that player 1 must choose a move from D_t.
6. If $|D_t| = |\Delta_t| = 1$ then the game is solved.
7. If $|D_t| < |D_{t-1}|$ or $|\Delta_t| < |\Delta_{t-1}|$ let $t = t + 1$ and goto 2.
8. Otherwise, we have reduced the game to the simplest form we can by this method.

Example (Iterated Elimination of Dominated Strategies)

Consider a game with the following payoff matrix:

	L	C	R
T	(4,3)	(5,1)	(6,2)
M	(2,1)	(8,4)	(3,6)
B	(3,0)	(9,6)	(2,8)

Look first at player 2's strategies…

Example (Iterated Elimination of Dominated Strategies)

C is strictly dominated by R, leading to:

	L	R
T	(4,3)	(6,2)
M	(2,1)	(3,6)
B	(3,0)	(2,8)

Player 1 *knows* that player 2 won't play C…

Example (Iterated Elimination of Dominated Strategies)

Conditionally, both M and B are dominated by T:

	L	R
T	(4,3)	(6,2)

Player 2 *knows* that player 1 will play T and so, they play L.

Again, we have a deterministic "solution".

Purely Competitive Games

- In a purely competitive game, one players reward is improved only at the cost of the other player.
- This means, that if $R(d', \delta) = R(d, \delta) + x$ then $S(d', \delta) = S(d, \delta) - x$.
- Hence $R(d', \delta) + S(d', \delta) = R(d, \delta) + S(d, \delta)$.
- The sum over all players' rewards is the same for all sets of moves.
- It doesn't change the domination structure or the ordering of expected rewards if we add a constant to all rewards.
- Hence, any purely competitive game is equivalent to a game in which:

$$\forall \delta \in \Delta, d \in D : R(d, \delta) + S(d, \delta) = 0$$

a *zero-sum game*.

Payoff and Zero-Sum Games

- In a zero-sum game:

$$S(d_i, \delta_j) = -R(d_i, \delta_j)$$

- Hence, we need specify only one payoff.
- Payoff matrices may be simplified to specify only one reward[6]

Example (Rock-Paper-Scissors is a zero-sum game)

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

- It can be convenient to use standard matrix notation, with $M = (m_{ij})$ and $R(d_i, \delta_j) = m_{ij}$.

[6] In the two player case, at least.

What if no move is dominant?

- In the RPS game, like many others, no move is dominant (or dominated) for either player.
- If either player commits themself to playing a particular move, the other play can exploit that commitment (if they knew what it was, that is).
- We need a strategy for dealing with such games.
- Perhaps the maximin approach might be useful here...

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 minimises:

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

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

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

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

- Which leads to a payoff for player 1 of:

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

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?

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:

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

  - 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:

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

- 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.

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

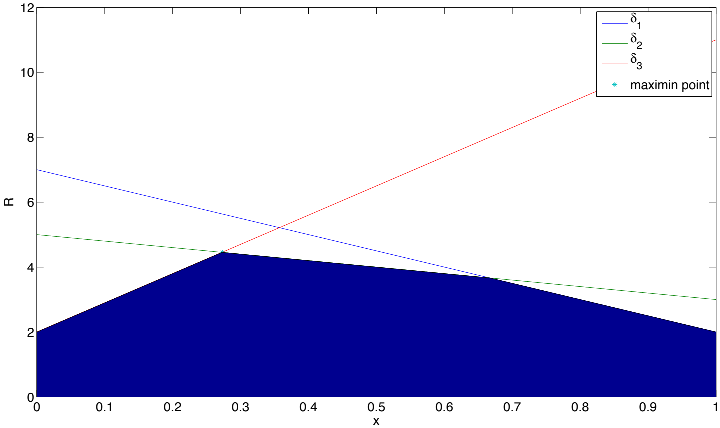

- The maximin response maximises the return in the worst case.
- In terms of our graph, this means we choose x to maximise the distance between the lowest of the lines and the ordinate axis.
- This is at the point where the lines associated with δ_2 and δ_3 intersect, at $x^\star$ which solves:

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

- 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:

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

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

$$\begin{aligned} V_2 = 3y + 11(1-y) =& 49/11 \\ 8y =& (121-49)/11 = 72/11 \Rightarrow y^\star = 9/11 \end{aligned}$$

- 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

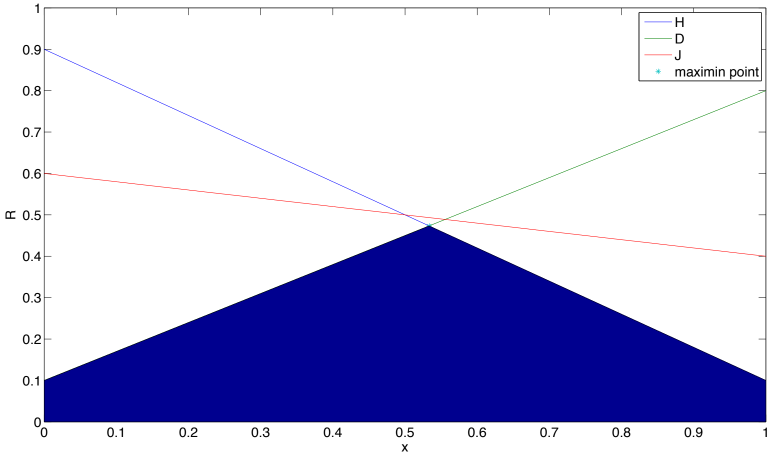

Example

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

$$p_H = p_D \Rightarrow 9 - 8x = 1 + 7x$$
$$8 = 15x \qquad \Rightarrow x^\star = 8/15$$

- 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.

A Few Useful Concepts from Game Theory

- Maximin pairs provide a “solution” concept for zero-sum games.
- Some problems arise considering non-zero-sum games:
  - Maximin pairs don’t necessarily make sense any more.
  - It’s not obvious what properties a solution should have.
- In general, we consider ideas of equilbrium and stability.
- Notions of optimality and equilibrium:
  - Pareto optimality.
  - Nash equilibrium.

Pareto Optimality

- A collection of strategies (one per player) in a game is *(strongly) Pareto optimal/efficient* if no change can be made which will improve one players reward without harming any other player.
- A collection of strategies is *weakly Pareto optimal* if no change can be made which will improve all players' rewards.
- If a collection of strategies is not Pareto optimal then at least one player could obtain a better outcome with a different collection.
- In a game of pure conflict, all sets of pure strategies are Pareto optimal.

Nash Equilibrium

- A collection of strategies (one per player) in a game is a *Nash equilibrium* if no player can improve their reward by unilaterally changing their strategy.
- In the two-player case, mixed strategies $\underline{x}$ and $\underline{y}$ comprise a Nash equilibrium if:

$$\forall \underline{x}' : \quad \bar{R}(\underline{x},\underline{y}) \geq \bar{R}(\underline{x}',\underline{y})$$
$$\forall \underline{y}' : \quad \bar{S}(\underline{x},\underline{y}) \geq \bar{S}(\underline{x},\underline{y}')$$

where

$$\bar{R}(\underline{x},\underline{y}) = \sum_{i=1}^{n}\sum_{j=1}^{m} x_i R(d_i,\delta_j) y_j \quad \bar{S}(\underline{x},\underline{y}) = \sum_{i=1}^{n}\sum_{j=1}^{m} x_i S(d_i,\delta_j) y_j$$

- If the inequality holds strictly we have a *strict Nash equilibrium*.

Nash Equilibria in 2 Player Zero Sum Games

- Maximin pairs are equivalent to Nash equilibria: if $\underline{x}^\star$ and $\underline{y}^\star$ are maximin, then, by definition:

$$\begin{aligned} \forall \underline{x}' : \quad & \bar{R}(\underline{x}^\star, \underline{y}^\star) \geq \bar{R}(\underline{x}', \underline{y}^\star) \\ \forall \underline{y}' : \quad & \bar{S}(\underline{x}^\star, \underline{y}^\star) \geq \bar{S}(\underline{x}^\star, \underline{y}') \end{aligned}$$

  A similar argument holds in the reverse direction.
- All equilibria have the same expected payoff (this follows from the fact that $S = -R$).
- These properties do not extend to non zero-sum games.

Nash Equilibria and the Prisoner's Dilemma

- Recall the prisoner's dilemma:

	S	B
S	(-1,-1)	(-5,0)
B	(0,-5)	(-4,-4)

- (B, B): both players betraying one another is a pure-strategy Nash equilibrium.
- (S, S): both players remaining silent is Pareto optimal: no change can be made which leads to improvement for one player and no worsening of the other player's situation.
- The (S, S) strategy set is not stable: it is not an equilibrium as either player can unilateral improve their own reward.

Solutions I: The Nash Sense

- Two pairs $(\underline{x}, \underline{y})$ and $(\underline{x}', \underline{y}')$ are interchangeable with respect to some property if $(\underline{x}', \underline{y})$ and $(\underline{x}, \underline{y}')$ have the same property.
- A game is *Nash solvable* if all equilibrium pairs are interchangeable (with respect to being equilibrium pairs).
- All zero-sum games are Nash solvable.
- Not many other games are.

Solutions II: The Strict Sense

- A game is *solvable in the strict sense* if:
  - Amongst the Pareto optimal pairs there is at least one equilibrium pair.
  - The equilibrium Pareto optimal pairs are interchangeable.
- The solution to such a game is the set of equilibrium Pareto optimal pairs.
- In a zero sum game, all strategies are Pareto optimal and so this reduces to the notion of Nash solvability: all zero sum games are solvable in the strict sense.

Solutions III: The Completely Weak Sense

- A game is *solvable in the completely weak sense* if after iterated elimination of dominated strategies, the reduced game is solvable in the strict sense.
- The solution is then the strict solution of the reduced game.
- In a zero sum game no strategies are dominated and so this reduces to the notion of solvability in the strict sense: all zero sum games are solvable in the completely weak sense.

Solutions and the Prisoner's Dilemma

- The only equilibrium pair of this game is (B, B).
- The only Pareto optimal strategy is (S, S).
- The game is Nash Solvable, with solution (B, B).
- The game is not solvable in the strict sense: no Pareto efficient pair of strategies is an equilibrium pair.
- The game is solvable in the completely weak sense:
  - S is a dominated strategy for both players.
  - The reduced game after IEDS has a single strategy (B) for each player.
  - The strategy (B, B) is Pareto efficient in the reduced game (no other strategy exists).
  - (B, B) is an equilibrium pair in the reduced game.
  - The solution set is (B, B).