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

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

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:

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

- Whilst, for player 2, we have

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

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.