12 Game Theory

12.1 What Is a Game?

- A **game** consists of: players, actions, and payoffs

  - A set of N players, indexed $i = 1, ...N$
  - A set of **actions** (or strategies) for each player: $a_i \in A_i$
    - ▷ An **action profile** characterizes the actions of all players

$$a = \{a_1, a_2, ...a_N\}$$

  - A mapping from action profiles to payoffs: $u_i(a)$ for all i

- Games are strategic: each player's payoff (potentially) depends on all players' actions

- Example: "I'm thinking of a number between 1 and 100"

  - Player 1 and Player 2, Bill and Ted, Bert and Ernie, whatever
  - Possible actions for each player:
    - ▷ Bert chooses any number from 1 to 10: $a_b \in \{1, 2, 3, ...100\}$
    - ▷ Ernie guesses any number from 1 to 10: $a_e \in \{1, 2, 3, ...100\}$
  - What's missing? Need to specify payoffs
  - Payoffs for each player:
    - ▷ Bert "wins" if Ernie doesn't guess his number (utility of winning?)
    - ▷ Ernie "wins" if he guesses the correct number (utility of winning?)

- This is what is missing from our model of equilibrium

  - Pretty clear what payoffs look like when supply equals demand
  - What are out-of-equilibrium payoffs?

12.2 Normal-Form Games

- In a **normal-form game**, players move simultaneously
  - Neither knows for certain what the other(s) will do
- An example of a 2×2 normal-form game:

		Player 2	
		left	right
Player 1	up	1,2	3,4
	down	5,6	7,8

  - 2 players (Player 1 and Player 2)
  - Player 1 chooses up or down
  - Player 2 chooses left or right
  - Resulting payoffs for each player (e.g. if P1 plays up and P2 plays left)
- **Example 1: Prisoner's Dilemma**

		Player 2	
		silent	talk
Player 1	silent	$-1,-1$	$-5,0$
	talk	$0,-5$	$-3,-3$

- Setup: two prisoners are arrested and kept in separate cells; they can keep silent, or they can talk to the police and help the authorities with their case
  - They are both better off if they keep silent, but...
- If Player 1 thinks Player 2 is going to keep silent, what will Player 1 do? (talk)
- If Player 1 thinks Player 2 is going to talk, what will Player 1 do? (talk)

- Player 1 prefers to talk, regardless of Player 2's action (circle)
  - The action "keep silent" is a **strictly dominated strategy**
- If Player 1 is trying to get the highest possible payoff, she will talk (circle it)
- *Repeat for Player 2*
- This suggests that both players will talk, even though they would earn higher payoffs if they could both agree to keep silent instead
  - Each player has an incentive to "free ride"
- **Example 2: Battle of the Sexes***
- AKA Bach or Stravinsky or Basketball or Skiing

		Morgan	
		basketball	skiing
Avery	basketball	2,1	0,0
	skiing	0,0	1,2

- If Morgan chooses basketball, what will Avery choose?
- If Morgan chooses skiing, what will Avery choose?
- In Battle of the Sexes, no action is strictly dominated
- Battle of the Sexes is an example of a **coordination game**; neither player has an incentive to free ride, and they are both better off if they are able to coordinate on one action

- **Example 3: Rock, Paper, Scissors**

	rock	paper	scissors
rock	0,0	−1,1	1,−1
paper	1,−1	0,0	−1,1
scissors	−1,1	1,−1	0,0

- Not all games are 2× 2...
- If Column chooses rock, what will Row do? (circle paper) If C chooses paper, what will R do? (circle scissors) If C chooses scissors, what will R do? (circle rock)
- What about C? If R chooses rock, what will C do? (circle paper) If R chooses paper, what will C do? (circle scissors) If R chooses scissors, what will C do? (circle rock)
- I haven't circled the same outcome/payoffs twice ⇒ it is hard to predict what will happen in this game (consistent with our experience of playing it)
- Rock, Paper, Scissors is an example of a **zero-sum game**
- **Example 4: A Public Goods Game**
- N players indexed $i = 1, 2, ..., N$
- Each player receives an endowment e_i and divides it between:
  - A pooled account that is divided equally (contributes c_i)
  - A private account that only benefits her ($e_i - c_i$)
- Player i's payoff: $e_i - c_i + \frac{1}{N}\sum_i \lambda c_i$ for some $1 < \lambda < N$
- Notice: all players are better off if everyone contributes everything, but **individually** each player has an incentive to free ride (just as in the PD game)

- **Summary**
- A game: players, actions, payoffs
- A normal form game: simultaneous moves
- A 2×2 normal form game: 2 players (moving simultaneously), each has two possible actions
- Basic game structures: coordination games, free riding/public goods games, zero-sum games

12.3 Iterated Elimination of Strictly Dominated Strategies

- An action is **strictly dominated** if a player always strictly prefers something else: irrespective of the other players' actions, we can rule it out

		Player 2		
		left	center	right
Player 1	up	1,0	1,2	0,1
	down	0,3	0,1	2,0

- In the game above, what would P1 do if she knew P2 were going to play left? (circle up) What would P1 do if she knew P2 were going to play center? (circle up) What would P1 do if she knew P2 were going to play right? (circle down)
  - So, neither up nor down is strictly dominated for P1
- What about P2? If P2 thought P1 was going to play up, what would P2 do? (circle center) If P2 thought P1 was going to play down, what would P2 do? (circle left)
  - Playing right is never a best response for P2
  - P1 might therefore expect that she will never choose it (cross it out)
- Now, if P1 "knows" P2 is not going to play right, down becomes a strictly dominated strategy in the game-within-a-game that is left (cross out down)

- Given this, P2 ought to realize that P1 will never play down (since P1 knows P2 will never play right), so left is also a strictly dominated strategy in the remaining game-within-in-game
- This process is called the **iterated elimination of strictly dominated strategies**, and sometimes it can lead us to a prediction about the likely outcome of a game... though other times it can't
  - Draw example of Battle of the Sexes
  - Show that no strategy is strictly dominated

12.4 Nash Equilibrium

- John Nash: won the Nobel Prize in 1994
- Proposed the idea that chosen actions in a game are "best responses"

An action profile $a^* = \{a_1^*, ..., a_N^*\}$ is a **Nash equilibrium** if every player's action is a best response to all the other players' actions, i.e. for every player i, the action a_i^* is a best response to the actions $a_1^*, ..., a_{i-1}^*, a_{i+1}^*, ..., a_N^*$

What is a **best response**?
When the other players' actions are $a_1^*, ..., a_{i-1}^*, a_{i+1}^*, ..., a_N^*$,
there's no (other) action $a_i \neq a_i^*$ that generates strictly higher utility than a_i^*

- Nash equilibria in Battle of the Sexes:

		Morgan	
		basketball	skiing
Avery	basketball	2,1	0,0
	skiing	0,0	1,2

- If Avery chooses basketball, Morgan's best response is basketball (circle); if Morgan chooses basketball, Avery's best response is basketball (circle)
- Both $\{b, b\}$ and $\{s, s\}$ are Nash equilibria; $\{b, s\}$ and $\{s, b\}$ are not
- Nash equilibrium in the Prisoner's Dilemma:
- Talk (defect) is the unique best response
- The only Nash equilibrium: $\{t, t\}$
- General points about Nash equilibria in 2×2 normal-form games:

Player 2

Player 1		left	right
	up	a,w	b,x
	down	c,y	d,z

- $\{U, L\}$ is a Nash equilibrium whenever $a \geq c$ and $w \geq x$
  - Notice that the inequalities are not strict
  - An action must be a best response, not the unique best response
- Assuming there are no "ties" (i.e. equal payoffs), what is the maximum number of (pure strategy) Nash equilibria that are possible?
- When will no (pure strategy) Nash equilibrium exist?

12.4.1 Duopoly: Cournot (1838)

- We can also characterize Nash equilibrium in games with continuous action spaces (as in the public goods game example discussed earlier)
- An economy contains two firms: Firm 1 and Firm 2
- Economy also contains consumers, represented by a demand function

- Total output: $Q = q_1 + q_2$
- Inverse demand: $P(Q) = a - Q = a - (q_1 + q_2)$
  - As in monopoly, each firm's choice of q_i impacts equilibrium price, so $MR \neq p$
- Firm i's costs: $C_i(q_i) = cq_i$
- What is each profit-maximizing firm's best response? Firm i wants to maximize:

$$\begin{aligned}
\pi_i\left(q_i|q_j\right) &= R_i\left(q_i|q_j\right) - C_i(q_i) \\
&= q_i\left[P(q_i + q_j)\right] - C_i(q_i) \\
&= q_i\left[a - (q_i + q_j)\right] - cq_i \\
&= aq_i - q_iq_j - cq_i - q_i^2
\end{aligned}$$

- Each firm's best response:

$$\begin{aligned}
\partial\pi_i/\partial q_i &= 0 \\
\Leftrightarrow a - q_j - c &= 2q_i \\
\Leftrightarrow q_i &= \left(a - c - q_j\right)/2
\end{aligned}$$

- Firm i does not **know** q_j, but in equilibrium it must be the case that I am best-responding to what I expect the other player to do
- This must characterize both Firm 1's best response and Firm 2's best response, so

$$q_1 = \left(a - c - q_2\right)/2$$

and

$$q_2 = \left(a - c - q_1\right)/2$$

- Plugging the second equation into the first:

$$\begin{aligned}
q_1 &= (a-c-q_2)/2 \\
\Leftrightarrow q_1 &= \left(a - c - \underbrace{\left(\frac{a-c-q_1}{2}\right)}_{=q_2} \right)/2 \\
\Leftrightarrow 2q_1 &= a - c - \frac{a-c-q_1}{2} \\
\Leftrightarrow 4q_1 &= 2a - 2c - a + c + q_1 \\
\Leftrightarrow 3q_1 &= a - c \\
\Leftrightarrow q_1^* &= (a-c)/3 \\
&= q_2^*
\end{aligned}$$

- Total quantity supplied is $\frac{2(a-c)}{3}$
- This is greater than what we'd expect in a monopoly: $q^M = \frac{a-c}{2}$
- Also lower than what we'd expect under perfect competition: $q^{PC} = a - c$
- We can also look at the problem graphically:

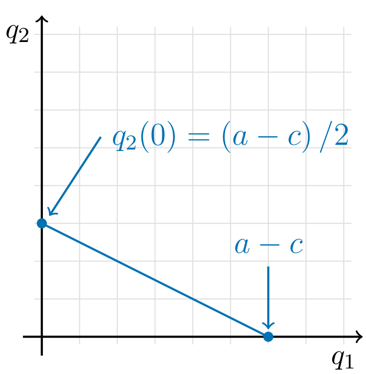

Firm i's best response:

$$q_i^*(q_j) = \frac{a-c-q_j}{2}$$

Best response is linear in q_j:

$$q_i^*(0) = \frac{a-c}{2}$$

$$q_i^* = 0 \Leftrightarrow q_j^* = a - c$$

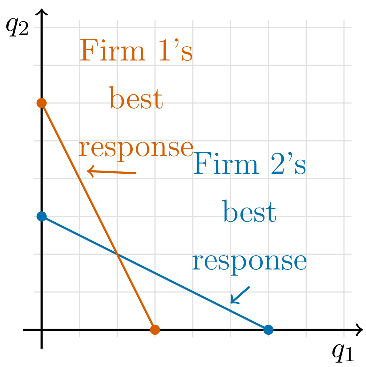

Firm i's best response:

$$q_i^*\left(q_j\right) = \frac{a-c-q_j}{2}$$

Best response is linear:

$$q_i^*\left(0\right) = \frac{a-c}{2}$$

$$q_i^* = 0 \Leftrightarrow q_j^* = a-c$$

12.5 Mixed Strategies

- Not every game has a (pure strategy) Nash equilibrium (e.g. Rock Paper Scissors)

A **pure strategy**: playing a single action with probability one

A **mixed strategy**: playing multiple actions with positive probability

$\Rightarrow$ A mixed strategy involves randomizing across actions

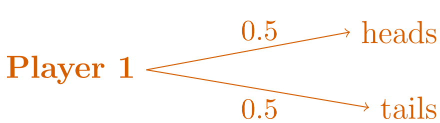

	heads	tails
heads	-1,1	1,-1
tails	1,-1	-1,1

- Imagine flipping a coin immediately prior to playing the game, so that your action is a random variable

- The other player doesn't know your action, only the probability of each possible action

- A **mixed strategy Nash equilibrium** is one where at least one player uses a mixed strategy

		Player 2	
		r heads	$1-r$ tails
Player 1	q heads	$-1,1$	$1,-1$
	$1-q$ tails	$1,-1$	$-1,1$

- What are the probabilities of the four possible outcomes of Matching Pennies?
- When does it make sense for a Player to use a mixed strategy?
  - What is P1's expected payoff from using a mixed strategy?

$$
\begin{aligned}
&= qr(-1) + q(1-r)(1) + (1-q)r(1) + (1-q)(1-r)(-1) \\
&= -qr + q - qr + r - qr - 1 + q + r - qr \\
&= 2r - 4qr - 1 + 2q \\
&= 2r(1-2q) - (1-2q) \\
&= (2r-1)(1-2q)
\end{aligned}
$$

  - When $r > 1/2$, $2r - 1 > 0$, so P1's expected payoff is decreasing in $q \Rightarrow$ P1 wants to make q as small as possible, so her best response is to play tails with probability one
  - When $r < 1/2$, $2r - 1 < 0$, so P1's expected payoff is increasing in $q \Rightarrow$ P1 wants to make q as large as possible, so her best response is to play **heads** with probability one
  - When is choosing $q \in (0, 1)$ a best response for P1? Only when $r = 1/2$.
  - Whether P1's best response is a mixed strategy depends on whether P2 is using a strategy that makes P1 precisely indifferent between her two pure strategies.
  - Vice versa for P2

- In a 2 player, normal-form game, the mixed strategies (q^*, r^*) are a **Nash equilibrium** is q^* if a best response to r^* and vice versa
- Matching Pennies has 0 pure strategy Nash equilibria, but 1 mixed strategy Nash equilibrium
- Let's look at the coordination game **stag hunt**

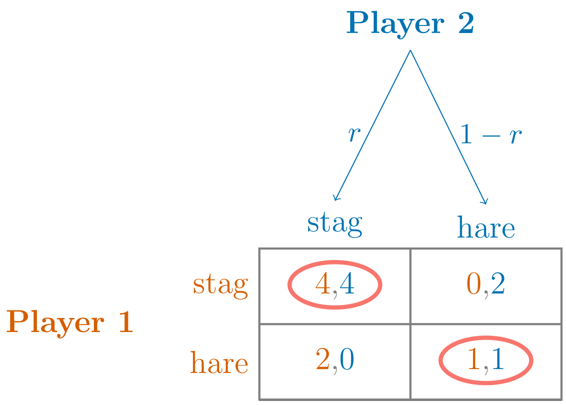

When is Player 1 indifferent?

$$\begin{aligned} u_1(\mathsf{stag}) &= u_1(\mathsf{hare}) \\ 4r &= 2r + (1-r) \\ r^* &= 1/3 \end{aligned}$$

- There are two pure strategy Nash equilibria in stag hunt
  - If P1 thinks P2 will play stag, P1 plays stag
  - It P1 thinks P2 will play hare, P1 plays hare
  - When is P1 indifferent between playing stag and playing hare?
- We've seen two cases:
  - In matching pennies, there is **no** pure strategy Nash equilibrium; there is only the one mixed strategy Nash equilibrium
  - In Stag Hunt, there are two pure strategy Nash equilibrium, and one mixed strategy Nash equilibrium
- How many mixed strategy Nash equilibria exist in Prisoner's Dilemma?
- A 2 player, normal-form game has either (with no payoff "ties"):
  - 1 pure strategy Nash equilibrium

- 1 mixed strategy Nash equilibrium
- 2 pure strategy equilibria and 1 mixed strategy Nash equilibrium

12.6 Dynamic Games

- In **dynamic games**, players move one after another (see below for structure)

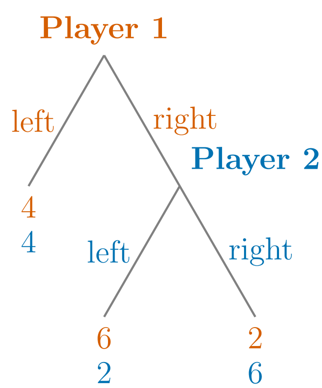

- *Translate game into players, actions, payoffs after walking through parts*
- Dynamic games also involve timing of moves
  - What is a decision node?
  - A strategy says what a player will do at all decision nodes
- Player 1 needs to form beliefs about what Player 2 will do (if Player 1 chooses R)
- In **backwards induction**, players divide the game into **subgames** and work backwards from the final subgame(s) to form beliefs about what other players will do
  - A **subgames** is a part of the game tree that starts from a decision node after the first decision node, containing all the branches that follow from the starting node of the subgame
  - Backwards induction relies on common knowledge of rationality
  - A backwards induction outcome is the outcome that results if we use backwards induction, in other words if all players assume all other players will always choose the higher payoff and expect other players to do the same (common knowledge)
  - A strategy profile is a **subgame perfect Nash equilibrium** if the actions chosen within each subgame are a Nash equilibrium (consistent with backwards induction)

- How would we solve Sequential BoS using backwards induction?
- **Stackelberg duopoly: duopoly as a dynamic game**

  Total output: $Q = q_1 + q_2$

  Inverse demand: $P(Q) = a - Q = a - (q_1 + q_2)$

  Firm i's costs: $C_i(q_i) = cq_i$

  Firm 1 chooses q_i before Firm 2

- Backwards induction $\rightarrow$ how will Firm 2 respond to Firm 1 (draw a timeline)?
- Firm 2 is going to maximize:

$$\begin{aligned}\pi_2\left(q_2|q_1\right) &= R_2\left(q_2|q_1\right) - C_2(q_2) \\ &= q_2\left[P(q_1+q_2)\right] - C_2(q_2) \\ &= q_2\left[a-(q_1+q_2)\right] - cq_2 \\ &= aq_2 - q_1q_2 - q_2^2 - cq_2 \\ &= \left(a-q_1-c\right)q_2 - q_2^2\end{aligned}$$

- Solving Firm 2's profit-maximization problem:

$$\begin{aligned}\partial\pi/\partial q_2 &= 0 \\ \Leftrightarrow a - q_1 - c - 2q_2 &= 0 \\ \Leftrightarrow q_2^* &= \left(a-q_1-c\right)/2\end{aligned}$$

- Knowing this, Firm 1 wants to maximize:

$$\begin{aligned}
\pi_1(q_1|q_2^*) &= R_1(q_1|q_2^*) - C_1(q_1) \\
&= q_1\left[P(q_1+q_2^*)\right] - C_1(q_1) \\
&= q_1\left[a - (q_1+q_2^*)\right] - cq_1 \\
&= q_1\left[a - \left(q_1 + \left(\frac{a-q_1-c}{2}\right)\right)\right] - cq_1 \\
&= q_1\left[a - q_1 - \frac{a}{2} + \frac{q_1}{2} + \frac{c}{2}\right] - cq_1 \\
&= q_1\left[\frac{a}{2} - \frac{q_1}{2} + \frac{c}{2}\right] - cq_1 \\
&= \left(\frac{a}{2} - \frac{c}{2}\right) q_1 - \frac{q_1^2}{2}
\end{aligned}$$

- Solving Firm 1's profit-maximization problem:

$$\begin{aligned}
\partial\pi/\partial q_1 &= 0 \\
\Leftrightarrow (a-c)/2 &= q_1^*
\end{aligned}$$

- Solving for Firm 2's best response:

$$\begin{aligned}
q_2^* &= \left(a - q_1^* - c\right)/2 \\
&= \left(a - \frac{a-c}{2} - c\right)/2 \\
&= (a-c)/4
\end{aligned}$$

- How does this compare to monopoly? Or perfect competition?

12.7 Behavioral Game Theory

- We are often just as interested in things (about human behavior) that game theory gets wrong as we are in the things game theory gets right

  Why is it called a Beauty Contest Game and not a Guessing Game?

  ⇒ Blame John Maynard Keynes

  Iterated elimination of dominated strategies ⇒ equilibrium near zero

  ⇒ Not actually a winning strategy much of the time

  ⇒ Some players may not arrive at the "optimal" strategy

  Rationality vs. common knowledge of rationality ("level-k thinking")

- Consider the Mini-Ultimatum Game:

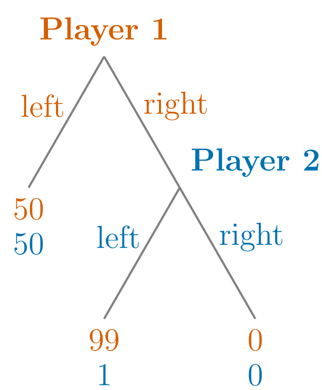

- Is the sub-game perfect Nash equilibrium prediction realistic?

Revised: November 30, 2022