Game Theory Basics

- Game theory is designed to model
  - How rational (payoff-maximizing) ``agents” will behave
  - When individual outcomes are determined by collective behavior.
  - Rules of a game specify agent payoffs as a function of actions taken by different agents.

Let’s play the median game

- On the index card, write down
  - Your name
  - An integer between 0 and 100 (inclusive).

- After we collect all the index cards, the person (or people) whose selected number is closest to 2/3 of the median of all the numbers (rounded down) wins a prize.

- E.g., if the numbers are 3, 4, 5, 38, 60, 70, 70, 90, 100

Prisoner's Dilemma

Prisoner II

Prisoner I

	stay silent	confess/betray
stay silent	$(-1, -1)$	$(-10, 0)$
confess/betray	$(0, -10)$	$(-8, -8)$

Confess/betray is best response no matter what.
dominates stay silent.

(Betray, Betray) is a dominant strategy.

Not a Pareto optimal strategy pair.

<u>Game</u>: # of players n.

for each player, S_i set of actions that player i can take.

$\forall$ strategy profile $\vec{s} = (s_1, s_2, .., s_n)$ $\quad (s_i \in S_i \ \forall i)$

$u_i(\vec{s})$ is payoff to player i when the players play strategy profile $\vec{s}$

$\vec{s}$ represents set of actions/strategies selected

$\in S_i$

Definition 2.2. A strategy s^* for player i is a **best response** to the strategies s_{-i} of others if it maximizes i's utility/payoff. That is

S_{-i} : strategy profile for all players but i

$$u_i(s^*, s_{-i}) \geq u_i(s, s_{-i})$$

for all $s \in S_i$.

Definition 2.3. Strategy s_i (strictly) dominates strategy s_i' if no matter what other players are doing, i's payoff playing s_i is at least as good (strictly better) than i's payoff playing s_i'.

$>$ ← strictly

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \quad \forall s_{-i} \forall s_i' \in S_i \setminus s_i.$$

If strategy s_i (strictly) dominates all strategies in S_i, then it is a (strictly) dominant strategy.

Definition 2.4. A strategy profile $(s_1, \ldots, s_n)$ is a **dominant strategy equilibrium** if for each player i, s_i is a dominant strategy.

	stay silent	confess/betray
stay silent	$(-1,-1)$	$(-10,0)$
confess/betray	$(0,-10)$	$(-8,-8)$

ISP 1 $s_1 \to t_1$

ISP 2 $s_2 \to t_2$

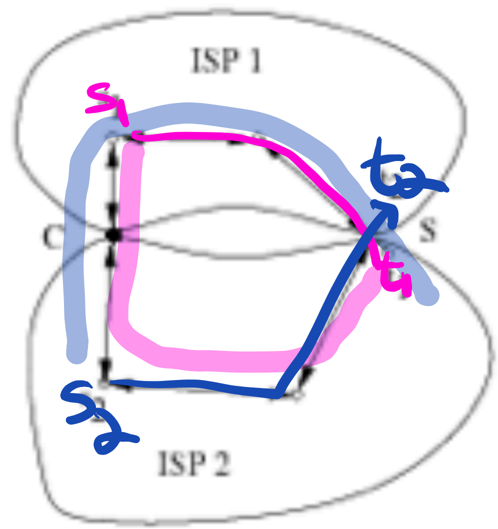

Another setting   P2P networks
- free riding

they each have a file that is desired by other.

decision: to upload desired file or not.

B

A

	upload	don't upload
upload	$(2,2)$	$(-1,3)$
don't upload	$(3,-1)$	$(0,0)$

benefit of receiving file   3
cost of uploading file   1

not uploading is a dominant strategy

Pollution Game: n countries, yes or no to legislation to control pollution / emissions

Pollution control costs 3
each country that pollutes adds 1 unit of cost to all countries

k countries are polluting $n-k-1$ aren't.

pollute. $k+1$

don't pollute $k+3$

dominant strategy to pollute.

Startup Game : Q: whether to enter a certain market or not.

Startup

Microsoft	Enter	Stay out
Enter	(2, −2)	(4, 0)
Stay out	~~(0, 4)~~	~~(0, 0)~~

for Microsoft
Entering dominates Staying out

Therefore Startup can safely assume that Microsoft will do so

$\underset{\text{(Microsoft)}}{\text{(Enter,}}$ $\underset{\text{(Startup)}}{\text{Stay Out)}}$ is a <u>Nash equilibrium</u>

i.e. each player is best responding to the other player i plays $s_i \in S_i$

Definition 2.8. A strategy profile $(s_1, \ldots, s_n)$ is a Nash equilibrium if for every i, s_i is a best response to s_{-i}.

Iterated Deletion of Dominated Strategies

b is weakly dominated by a

if

- Deletion of a dominated strategy: find a player i and a strategy $b \in S_i$ such that $a \in S_i$ weakly dominates strategy b. Delete strategy b from S_i.
- Update definition of what's dominated (assuming b will never be played).
- Iterate until no weakly dominated strategy remains.

$\forall s_{-i}$

$u_i(a, s_{-i}) \geq u_i(b, s_{-i})$

and $\exists \tilde{s}_{-i}$

where $u_i(a, \tilde{s}_{-i}) > u_i(b, \tilde{s}_{-i})$

b is strongly dominated by a if $\forall s_{-i}$

$u_i(a, s_{-i}) > u_i(b, s_{-i})$

If each player has only a single remaining strategy, we say that the game is *solvable* by iterated deletion of dominated strategies, and we say that iterated deletion of dominated strategies predicts that each player will play their only remaining response.

Caveat: only good prediction if deleted only strongly dominated strategies

otherwise prediction not unique (depends on order)

Back to the median game

- On the index card, write down
  - Your name
  - An integer between 0 and 100 (inclusive).
- After we collect all the index cards, the person (or people) whose selected number is closest to 2/3 of the median of all the numbers (rounded down) wins a prize.

most median

$\frac{2}{3}$ med $= 66$

Iterated Deletion of Dominated Strategies

- Deletion of a dominated strategy: find a player i and a strategy $b \in S_i$ such that $a \in S_i$ weakly dominates strategy b. Delete strategy b from S_i.
- Update definition of what's dominated (assuming b will never be played).
- Iterate until no weakly dominated strategy remains.

If each player has only a single remaining strategy, we say that the game is *solvable* by iterated deletion of dominated strategies, and we say that iterated deletion of dominated strategies predicts that each player will play their only remaining response.

Coordination Game

Bob.

Alice

	party	stay home
party	$2, 2$	$-1, 0$
stay home	$0, -1$	$0, 0$

Definition 2.2. A strategy s^* for player i is a **best response** to the strategies s_{-i} of others if it maximizes i's utility/payoff. That is

$$u_i(s^*, s_{-i}) \geq u_i(s, s_{-i})$$

for all $s \in S_i$.

Definition 2.8. A strategy profile $(s_1, \ldots, s_n)$ is a Nash equilibrium if for every i, s_i is a best response to s_{-i}.

Network coordination games

each node is person.

action set: use the app or not.

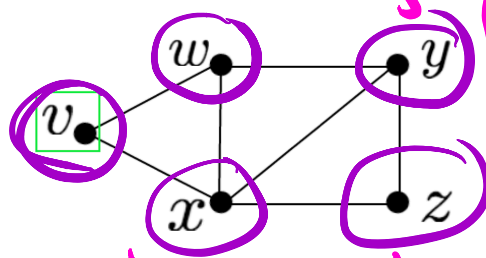

$$\text{payoff to } v = \begin{cases} 2 & \geq k_v \text{ of her friends use it \& she does} \\ -1 & < k_v \text{ of her \& she does} \\ 0 & \text{if she doesn't use it.} \end{cases}$$

NE

- 0's all around.
- all use it.

network cascade

Prog Proj 1:

$N, S, T, W, A_s, (\varphi_u)_{u \in U}\ (S_{s,u})_{u \in U}$

max aptitude

max quality

max synergy

usnews – highest quality

synergist – highest synergy

holist – highest quality + synergy

random

Parking Game.

inspector

	don't inspect ($\frac{9}{10}$)	inspect ($\frac{1}{10}$)
parker: legal ($\frac{4}{5}$)	(0, 0)	(0, −1)
parker: illegal ($\frac{1}{5}$)	(10, −10)	(−90, −6)

	dont inspect	inspect.
p	0, 0	0, −1
$1-p$	10, −10	−90, −6
	$-10(1-p)$	$-p-6(1-p)$
	$p > \frac{4}{5}$	$p < \frac{4}{5}$

$$-10(1-p) \geq -p-6(1-p)$$
$$p \geq 4(1-p)$$
$$5p \geq 4$$
$$p \geq \frac{4}{5}$$

	q	$1-q$	
legal	0, 0	0, −1	0
illegal	10, −10	−90, −6	$10q-90(1-q)$

legal better than illegal

$$0 \geq 10q-90(1-q)$$
$$90 \geq 100q$$
$$\frac{9}{10} \geq q$$

$\left(\frac{4}{5}, \frac{9}{10}\right)$ is a Nash Eq in mixed strategies

Parker: al with prob $\frac{4}{5}$ al with prob $\frac{1}{5}$ inspector inspects w/ $\frac{9}{10}$ doesn't $\frac{1}{10}$

$x_i(s)$ = Prob that player i plays strategy s.

Definition 4.1. A **mixed strategy** is a probability distribution over pure strategies.

Definition 4.2. A strategy profile $(x_1, \ldots, x_n)$ where each x_i is possibly a mixed strategy $x_i : S_i \to$ probability distribution and $\sum_{s \in S_i} x_i(s) = 1$ is a (mixed) Nash equilibrium if for each i,

$$\sum_{s \in S_i, s_{-i} \in S_{-i}} x_i(s) \operatorname{Prob}(s_{-i}) u_i(s, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \operatorname{Prob}(s_{-i}) u_i(s_i, s_{-i}).$$

Expected payoff higher using mixed strategy x_i than using any pure strategy when others play x_j (mixed strategy of player j)

Expected payoff / utility of player i when he plays s_i

	party	stay home
party	$2, 2$	$-1, 0$
stay home	$0, -1$	$0, 0$

This game has 2 pure Nash equilibria

It also has a mixed NE where each player parties with prob $\frac{1}{2}$ & stays home w/ prob $\frac{1}{2}$

$$\begin{array}{rcc} & p & 1-p \\ 3p-1 = 2p-(1-p) & 2,2 & -1,0 \\ 1-p & 0,-1 & 1,1 \end{array}$$

$3p-1 = 1-p$

$4p = 2$ $\quad p = \frac{1}{2}$ $\quad$ exp payoff $= \frac{1}{2}$. lower than both pure eq

Fact 4.4. **Indifference principle:** If a mixed strategy of player i in a Nash equilibrium randomizes over a set of pure strategies $T_i \subset S_i$, then the expected payoff to the player from each pure strategy in T_i must be the same. And the payoff from any other strategy must be at most this high.

Summary so far

- A Nash equilibrium is a set of stable (possibly mixed) strategies.
- Stable means that no player has an incentive to deviate given what the other players are doing.
- Pure equilibrium: there may be none, unique or multiple. Can be identified with “best response diagrams”.
- A joint mixed strategy for n players:
  - A probability distribution for each player (possibly different)
- It is an equilibrium if
  - For each player, their distribution is a best response to the others.
  - Only consider unilateral deviations.
  - Everyone knows all the distributions (but not the outcomes of the coin flips).
- Nash’s famous theorem: every game has a mixed strategy equilibrium.

Issues

- Does not suggest how players might choose between different equilibria
- Does not suggest how players might learn to play equilibrium.
- Does not allow for bargains, side payments, threats, collusions, “pre-play” communication.

- Computing Nash equilibria for large games is computationally difficult.

Other issues

- Relies on assumptions that might be violated in the real world
  - Rationality is common knowledge.
  - Agents are computationally unbounded.
  - Agents have full information about other players, payoffs, etc.

Zero-sum games

Penalty Kicks.

Goalie.

\#'s are payoff/gain of row player.

Kicker

		L	R
p	L	0.5	1
$1-p$	R	0.9	0.8

$0.5p + 0.9(1-p) = 0.9 - 0.4p$

$p + 0.8(1-p)$

$0.8 + 0.2p$

Suppose Kicker says L with prob p.

loss of Goalie if goes R

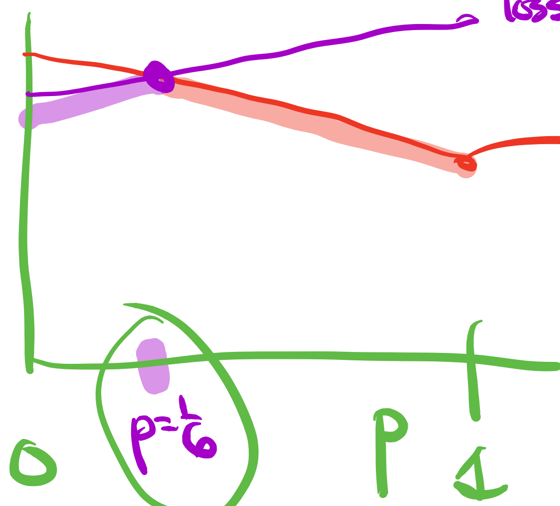

loss if Goalie goes L

$0.9 - 0.4p$.

$\max_p \min(\text{gain } G \to L, \text{gain } G \to R)$

"Kicker goes first"

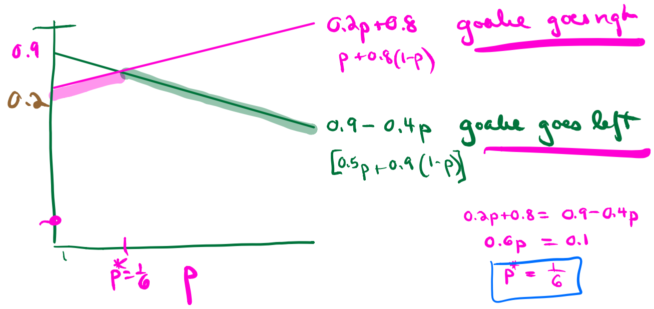

$0.2p+0.8 = 0.9-0.4p$

$0.6p = 0.1$

$\boxed{p^* = \frac{1}{6}}$

⟸ Suppose Kicker must announce p ahead of time (go first). What is kicker's best choice for p?

Choosing $p^*=\frac{1}{6}$ (kick left w/ prob $\frac{1}{6}$, kick right w/ prob $\frac{5}{6}$) maximizes kicker's expected gain if she has to announce first, i.e. is the p that maximizes $\min\big(p+0.8(1-p),\ 0.5p+0.9(1-p)\big)$

If Kicker plays p^*, she guarantees herself an expected payoff of $\frac{1}{6}+0.8\left(1-\frac{1}{6}\right)=\frac{5}{6}=V_1$

"Goalie goes first"

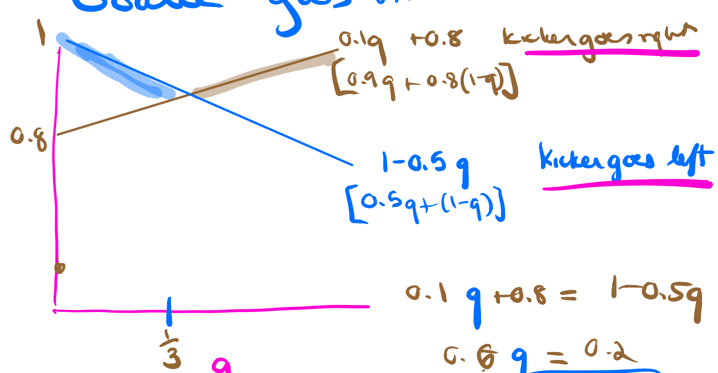

$0.1q+0.8 = 1-0.5q$

$0.6q = 0.2$

$\boxed{q^* = \frac{1}{3}}$

⟸ Suppose Kicker gets to best respond to Goalie's mixed strategy (q-dive left, 1-q, dive right) What is goalie's best choice for q?

$0.0\bar{3}$

Choosing $q^*=\frac{1}{3}$ (dive left with prob $\frac{1}{3}$ and dive right w/ prob $2/3$) minimizes goalie's expected loss if he has to announce first. i.e. minimizes $\max\big(0.9q+0.8(1-q),\ 0.5q+1-q\big)$

If Goalie goes first, he can guarantee himself a loss of at most $0.9\cdot\frac{1}{3}+0.8\frac{2}{3}=\frac{5}{6}=V_2$

V_1: exp payoff the kicker can guarantee himself if he has to go first.

V_2: exp payoff the kicker can guarantee himself if he can go second.

= min exp loss the goalie can guarantee if he has to go first.

$V_1 \leq V_2$

Theorem 4.5 (John von Neumann, 1928). *Let V_1 be the expected gain that player I (maximizer) can guarantee herself in the worst case, and let p^* be the mixed strategy that achieves V_1. Let V_2 be the lowest expected loss that player II (minimizer) can limit his loss to in the worst case and let q^* be the mixed strategy that achieves V_2.*

Then for any 2-player, zero sum game $V_1 = V_2 = V$ (called the minimax value of the game) and (p^*, q^*) is a Nash equilibrium.

$V_1 = V_2$

Summary – zero-sum games

- Zero-sum games have a "value".
- Optimal strategies are well-defined.
- Maximizer can guarantee a gain of at least V by playing p*
- Minimizer can guarantee a loss of at most V by playing q*.
- This is a Nash equilibrium.
- In contrast to general-sum games, optimal strategies in zero-sum games can be computed efficiently (using linear programming).

1500 penalty kicks

0.423 0.577 ← actual observed fractions

0.42 0.58 ← optimal strategies in game

			L	R
0.4	0.38	L	0.58	0.95
0.6	0.62	R	0.93	0.7

Extensive Form Games.

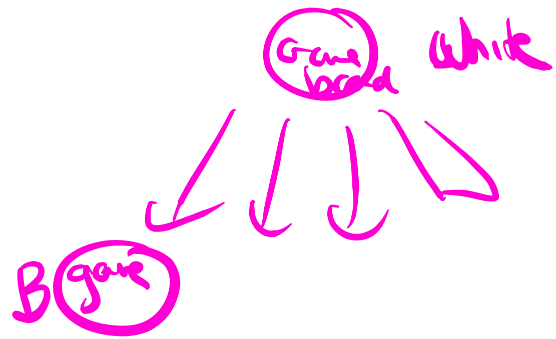

Definition 5.1. A k-player finite **extensive-form game** is defined by a finite, rooted tree T.

- Each node in T represents a possible state in the game, with leaves representing terminal states.
- Each internal (nonleaf) node v in T is associated with one of the players, indicating that it is his turn to play if/when v is reached.
- The edges from an internal node to its children are labeled with **actions**, the possible moves the corresponding player can choose from when the game reaches that state.
- Each leaf/terminal state results in a certain payoff for each player.

Mutual Assured Destruction.

A aggressive, B benign

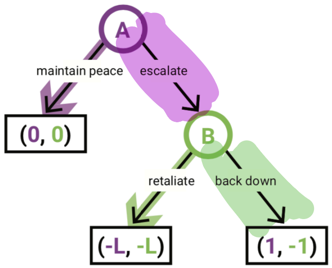

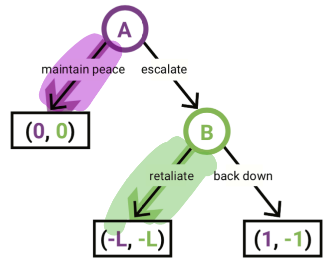

subgame perfect NE

2 NE

A **pure strategy** for a player in an extensive-form game specifies an action to be taken at each of that player's nodes.

A **mixed strategy** is a probability distribution over pure strategies.

The kind of equilibrium that is computed by backward induction is called a **subgame-perfect equilibrium** because the behavior in each **subgame**, is also an equilibrium.

Centipede,

FIGURE 6.4. The top part of the figure shows the game and the resulting payoffs at each leaf. At each node, the "greedy" strategy consists of following the downward arrow, and the "continue" strategy is represented by the arrow to the right. Backward induction from the node with pot-size 99 shows that at each step the player is better off being greedy.

so far games of perfect info

Large Company vs Startup

- startup announces a technology that threatens big company.
- $\frac{1}{2}$ is prob that BC can pull together comp product.

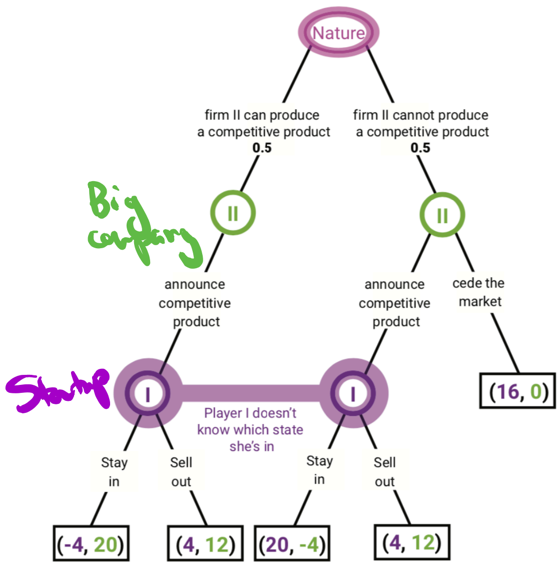

FIGURE 6.11. The figure shows the Large Company vs Startup game. Prior to the beginning of the game, player I announces her new technology. At the beginning of the game, there is a move by nature, which determines whether or not II actually can pull together a competitive product. Only player II is privy to the outcome of this move by nature. The two nodes at which player I makes a move form a single information set: player I does not know which of these states she is in. All she knows is that II has announced a competitive product, and knowing only that, she has to decide between competing with the giant or letting the giant buy her out. Thus, her strategy is the same at both nodes in the information set.

player I \ player II	announce/cede	announce/announce
stay in (I)	$(6, 10)$	$(8, 8)$
sell out (O)	$(10, 6)$	$(4, 12)$

Repeated Prisoners Dilemma.

		player II	
		cooperate (C)	defect (D)
player I	cooperate (C)	$(6, 6)$	$(0, 8)$
	defect (D)	$(8, 0)$	$(2, 2)$

Infinitely repeated game with discounting.

Fix $\beta < 1$

$$\text{discounted payoff} = \sum_{t=1}^{\infty} \beta^{t} \left(\text{the players payoff in round } t\right)$$

- β prob game continues.

Grim Trigger: Cooperate until a round in which your opponent defects, then defect from then on.

(Grim Trigger, Grim Trigger)
NE
(TitForTat, TitForTat)

		player II	
		cooperate (C)	defect (D)
player I	cooperate (C)	$(6,6)$	$(0,8)$
	defect (D)	$(8,0)$	$(2,2)$

					t	$t{+}1$	...	
Player I Grim Trigger	C	C	C	...	C	D	D	D
Player II deviates	C	C	C	...	D	D	D	D
Payoff vector	(6, 6)				(0, 8)	(2, 2)		
versus					(6, 6)	(6, 6)	...	
					with no deviation			

$$8\beta^t + 2\sum_{j>t}\beta^j \quad \text{vs} \quad 6\sum_{j\geq t}\beta^j$$

$$6\sum_{j\geq t}\beta^t > 8\beta^t + 2\sum_{j>t}\beta^j \quad \text{when} \quad 2\beta^t < 4\sum_{j>t}\beta^j = \frac{4\beta^{t+1}}{1-\beta}$$

$$1-\beta < 2\beta \Rightarrow \beta > \tfrac{1}{3}$$

Tit for Tat.

Cooperate in round 1
in every round $k > 1$
play what your opponent played in round $k-1$

player II

player I	cooperate (C)	defect (D)
cooperate (C)	$(6,6)$	$(0,8)$
defect (D)	$(8,0)$	$(2,2)$

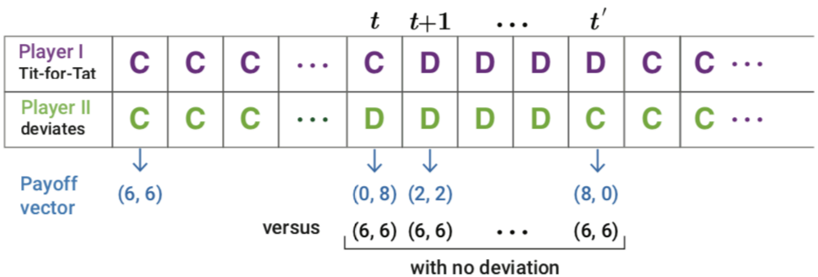

Figure 6.14. Illustration of deviation in Tit-for-Tat strategies