Lec 7: Equilibirium Notions

CS 580 (AGT)

Instructor: Ruta Mehta

Games

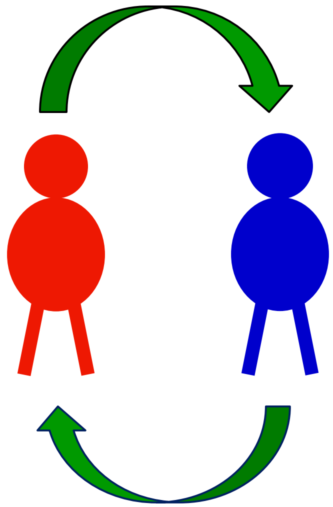

Payoffs

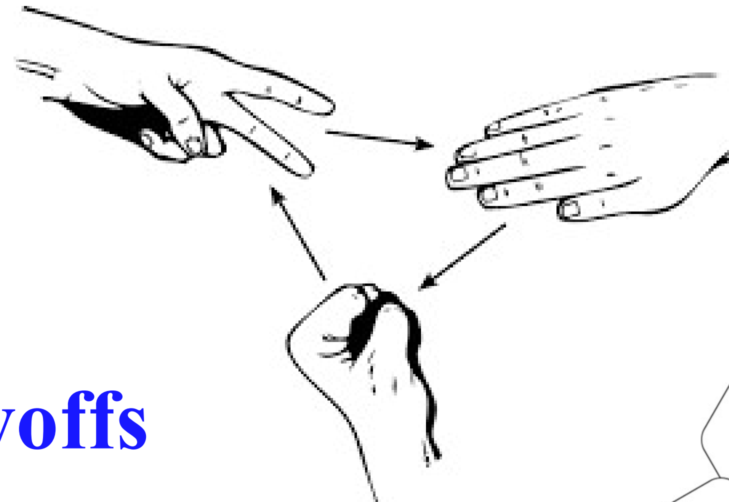

Players

Strategies

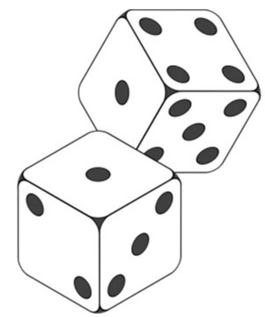

Randomize!

Games

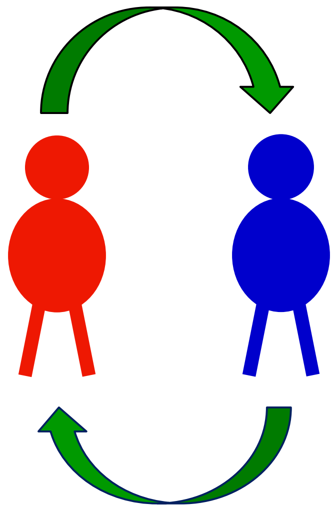

Players

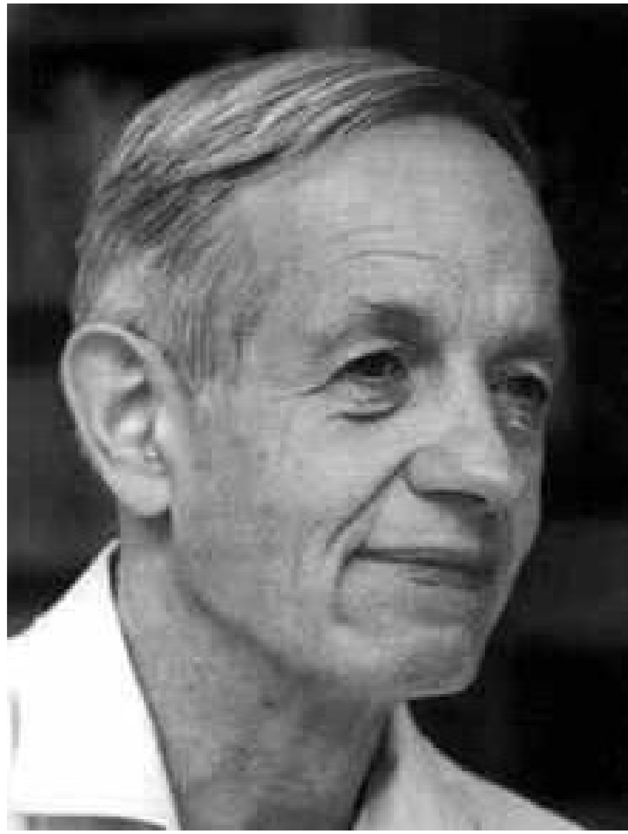

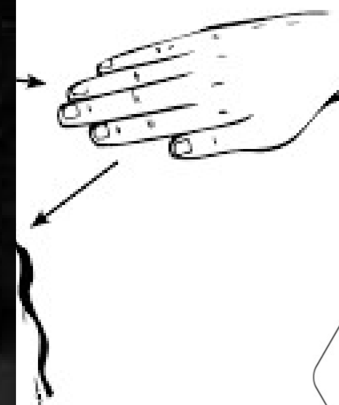

tegies

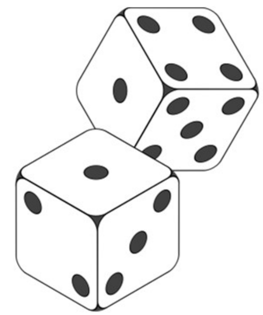

Randomize!

Nash (1950):

There exists a state where no player gains by unilateral deviation.

Nash equilibrium (NE)

Rock-Paper-Scissors

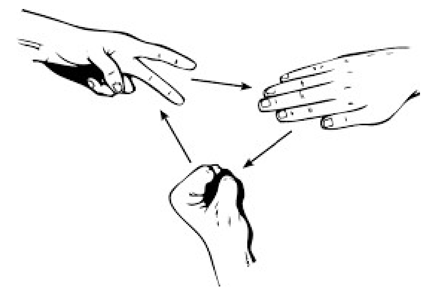

	R	P	S
R	0 0	-1 1	1 -1
P	1 -1	0 0	-1 1
S	-1 1	1 -1	0 0

No pure stable state!

Nash Eq.: No player gains by deviating individually

Both playing (1/3,1/3,1/3) is the only NE.

Why?

Formally: Games and Nash Equilibrium

- N: Set of players/agents $N = \{1, 2\}$
- $i \in N, S_i$: Set of strategies/moves of player i $S_1 = S_2 = \{R, P, S\}$
- $s = (s_1, \ldots, s_n) \in X_i S_i$, $u_i(s)$: payoff/utility of player i

  (R, P) $u_1(R, P) = -1$ $u_2(R, P) = 1$
- $\sigma_i \in \Delta(S_i)$ randomized strategy of i
- ***Nash equilibrium:*** $\sigma = (\sigma_1, \ldots, \sigma_n)$ *s.t.*

  $\forall i \in N, \quad u_i(\sigma_i, \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}), \quad \forall \tau_i \in \Delta(S_i)$

  $u_i(\sigma_i, \sigma_{-i}) = \max_{\tau_i} u_i(\tau_i, \sigma_{-i})$

  σ_{-i}: strategy of all except i

$u_1(R, \sigma_2) = -0.3$, $u_1(P, \sigma_2) = 0.5$ $u_1(S, \sigma_2) = 0.5$

Formally: Games and Nash Equilibrium

- N: Set of players/agents
- $i \in N, S_i$: Set of strategies/moves of player i
- $s = (s_1, \ldots, s_n) \in X_i S_i$, $u_i(s)$: payoff/utility of player i

- $\sigma_i \in \Delta(S_i)$ randomized strategy of i

- ***Nash equilibrium:*** $\sigma = (\sigma_1, \ldots, \sigma_n)$ *s.t.*
  $\forall i \in N, \quad u_i(\sigma_i, \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}), \quad \forall \tau_i \in \Delta(S_i)$

- **Observation:** a player randomizes only among those pure strategies that give her maximum payoff.

Prisoner's Dilemma

Two thieves caught for burglary.

Two options: {confess, not confess}

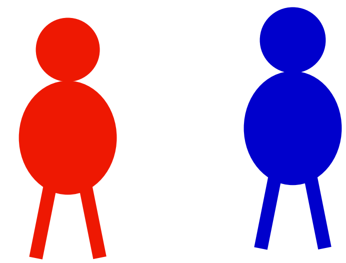

	N	C
N	-1 -1	-6 0
C	0 -6	-5 -5

Prisoner's Dilemma

Two thieves caught for burglary.

Two options: {confess, not confess}

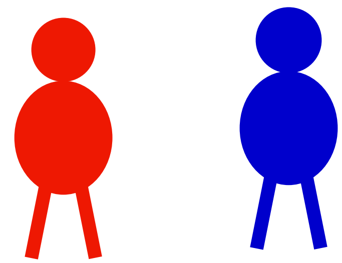

	N	C
N	-1 -1	-6 0
C	0 -6	-5 -5

Dominant Strategy Equilibrium (DSE)

Prisoner's Dilemma

Two thieves caught for burglary.

Two options: {confess, not confess}

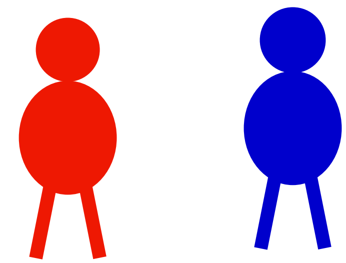

	N	C
N	-1 -1	-6 0
C	0 -6	-5 -5

DSE: $s = (s_1, \ldots, s_n)\ s.t.$

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

Incomplete Information: Bayesian

- Utility of a player depends on her **type** and the actions taken in the game
  - θ_i is player i's type, $\theta_i \sim \Theta_i$. Utilily when θ_i type and s play is $u_i(\theta_i, s)$
  - Each player knows/learns its own type, but only distribution of others (before choosing action)
    - Pure strategy $s_i : \Theta_i \rightarrow S_i$ (where S_i is i's set of actions)

(In general players can also receive signals about other players' utilities; we will not go into this)

row player (Alice)
type 1 (prob. 0.5)

	L	R
U	4	6
D	2	4

column player (Bob)
type 1 (prob. 0.5)

	L	R
U	4	6
D	4	6

row player
type 2 (prob. 0.5)

	L	R
U	2	4
D	4	2

column player
type 2 (prob. 0.5)

	L	R
U	2	2
D	4	2

Car Selling Game

- **A seller wants to sell a car**
- **A buyer has private value ‘v’ for the car w.p. P(v)**
- **Sellers knows P, but not v**
- **Seller sets a price ‘p’, and buyer decides to buy or not buy.**
- **If sell happens then the seller gets p, and buyer gets (v-p).**

S_1=All possible prices, Θ_1={1}
S_2={buy, not buy}, Θ_2 =**All possible ‘v’**
$U_1(1,(p,\text{buy}))=p,\quad U_1(1,(p,\text{not buy}))=0$
$U_2(v,(p,\text{buy}))=v-p,\quad U_2(v,(p,\text{not buy}))=0$

Bayes-Nash equilibrium

- A profile of strategies is a Bayes-Nash equilibrium if it is a Nash equilibrium for the normal form of the game
  - Mixed strategy of player i, $\sigma_i: \Theta_i \rightarrow \Delta(S_i)$
  - for every i, for every type θ_i, for every alternative action $z_i \in \Delta(S_i)$, we must have:

$$\Sigma_{\theta_{-i}} P(\theta_{-i})\, u_i(\theta_i, \sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})) \geq \Sigma_{\theta_{-i}} P(\theta_{-i})\, u_i(\theta_i, z_i, \sigma_{-i}(\theta_{-i}))$$

$$\Pi_{p \neq i} P(\theta_p)$$