Mixing Time and Stationary Expected Social Welfare of Logit Dynamics

Francesco Pasquale

Dipartimento di Informatica "R. Capocelli"
Università di Salerno

joint work-in-progress with

Vincenzo Auletta, Diodato Ferraioli, and Giuseppe Persiano

Roma, June 17, 2010

Outline

Framework Description

Examples

Research Directions

Game Theory

...in one slide

$$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$$

- $[n] = \{1, \ldots, n\}$ players;
- $\mathcal{S} = \{S_1, \ldots, S_n\}$; $\quad S_i = \{$ actions for player $i\}$;
- $\mathcal{U} = \{u_1, \ldots, u_n\}$; $\quad u_i : S_1 \times \cdots \times S_n \to \mathbb{R}$ utility functions

Game Theory

...in one slide

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$

- $[n] = \{1, \dots, n\}$ players;
- $\mathcal{S} = \{S_1, \dots, S_n\}$; $S_i = \{$ actions for player $i\}$;
- $\mathcal{U} = \{u_1, \dots, u_n\}$; $u_i : S_1 \times \cdots \times S_n \to \mathbb{R}$ utility functions

$\mathbf{x} = (x_1, \dots, x_n) \in S_1 \times \cdots \times S_n$ **pure Nash** equilibrium if for every $i \in [n]$ and for every $y \in S_i$

$$u_i(\mathbf{x}_{-i}, y) \leqslant u_i(\mathbf{x})$$

Game Theory

...in one slide

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$

- $[n] = \{1, \ldots, n\}$ players;
- $\mathcal{S} = \{S_1, \ldots, S_n\}$; $S_i = \{$ actions for player $i\}$;
- $\mathcal{U} = \{u_1, \ldots, u_n\}$; $u_i : S_1 \times \cdots \times S_n \to \mathbb{R}$ utility functions

$\mathbf{x} = (x_1, \ldots, x_n) \in S_1 \times \cdots \times S_n$ **pure Nash** equilibrium if for every $i \in [n]$ and for every $y \in S_i$

$$u_i(\mathbf{x}_{-i}, y) \leqslant u_i(\mathbf{x})$$

$\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n) \in \Delta(S_1) \times \cdots \times \Delta(S_n)$ **mixed Nash** equilibrium if for every $i \in [n]$ and for every $\sigma \in \Delta(S_i)$

$$\mathbf{E}_{(\boldsymbol{\mu}_{-i}, \sigma)}[u_i] \leqslant \mathbf{E}_{\boldsymbol{\mu}}[u_i]$$

Dynamics

Reaching equilibria

Dynamics: Choose a player, update her strategy, repeat.

Dynamics

Reaching equilibria

Dynamics: Choose a player, update her strategy, repeat.

Best response dynamics

At her turn, player i chooses the action $y \in S_i$ that maximizes her utility

$$u_i(\mathbf{x}_{-i}, y) \geqslant u_i(\mathbf{x}_{-i}, z) \text{ for every } z \in S_i$$

Dynamics

Reaching equilibria

Dynamics: Choose a player, update her strategy, repeat.

Best response dynamics

At her turn, player i chooses the action $y \in S_i$ that maximizes her utility

$$u_i(\mathbf{x}_{-i}, y) \geqslant u_i(\mathbf{x}_{-i}, z) \text{ for every } z \in S_i$$

Questions

- **Convergence**

Dynamics

Reaching equilibria

Dynamics: Choose a player, update her strategy, repeat.

Best response dynamics

At her turn, player i chooses the action $y \in S_i$ that maximizes her utility

$$u_i(\mathbf{x}_{-i}, y) \geqslant u_i(\mathbf{x}_{-i}, z) \text{ for every } z \in S_i$$

Questions

- Convergence

  If yes then...

- **Speed of convergence**.

Randomized Best Response

Logit Dynamics

Randomized Best Response: Choose for the next round strategy y with probability proportional to the returned utility.

Randomized Best Response

Logit Dynamics

Randomized Best Response: Choose for the next round strategy y with probability proportional to the returned utility.

$$\mathbf{x} \in S_1 \times \cdots \times S_n, \quad i \in [n], \quad y \in S_i \quad \Rightarrow \qquad \sigma_i(y \mid \mathbf{x}) \sim e^{\beta u_i(\mathbf{x}_{-i}, y)}$$

Randomized Best Response

Logit Dynamics

Randomized Best Response: Choose for the next round strategy y with probability proportional to the returned utility.

$$\mathbf{x} \in S_1 \times \cdots \times S_n, \quad i \in [n], \quad y \in S_i \quad \Rightarrow \qquad \sigma_i(y \mid \mathbf{x}) \sim e^{\beta u_i(\mathbf{x}_{-i}, y)}$$

$$\sigma_i(y \mid \mathbf{x}) = \frac{e^{\beta u_i(\mathbf{x}_{-i}, y)}}{\sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}}$$

$\beta =$ "Inverse noise"

Observation

- $\beta = 0$ players play uniformly at random;
- $\beta \to \infty$ players play best response (u.a.r. over best responses if more than one)

Logit Dynamics

Description

Logit dynamics [Blume, GEB'93]

From any profile $\mathbf{x}$, choose a player $i \in [n]$ u.a.r and update her action with probability $\sigma_i(\cdot \mid \mathbf{x})$.

$$\begin{tikzcd} \mathbf{x} \arrow[r, bend left, "p \sim e^{\beta u_i(\mathbf{x}_{-i}, y)}"] & (\mathbf{x}_{-i}, y) \end{tikzcd}$$

$$P\left(\mathbf{x}, (\mathbf{x}_{-i}, y)\right) = \frac{1}{n}\sigma_i(y \mid \mathbf{x})$$

Logit Dynamics

Description

Logit dynamics [Blume, GEB'93]

From any profile $\mathbf{x}$, choose a player $i \in [n]$ u.a.r and update her action with probability $\sigma_i(\cdot \mid \mathbf{x})$.

$$\mathbf{x} \xrightarrow{p \sim e^{\beta u_i(\mathbf{x}_{-i}, y)}} (\mathbf{x}_{-i}, y)$$

$$P\left(\mathbf{x}, (\mathbf{x}_{-i}, y)\right) = \frac{1}{n}\sigma_i(y \mid \mathbf{x})$$

This process defines an ergodic Markov chain

Markov chains

...in one slide

$$\{X_t \,:\, t \in \mathbb{N}\},\ (\Omega, P).$$

Markov chains

...in one slide

$\{X_t \,:\, t \in \mathbb{N}\}$, (Ω, P).

- **Irreducible**: for every $\mathbf{x}, \mathbf{y} \in \Omega$, $\exists t \in \mathbb{N} \,:\, P^t(\mathbf{x}, \mathbf{y}) > 0$;

Markov chains

...in one slide

$\{X_t \,:\, t \in \mathbb{N}\}$, (Ω, P).

- Irreducible: for every $\mathbf{x}, \mathbf{y} \in \Omega$, $\exists t \in \mathbb{N} \;:\; P^t(\mathbf{x}, \mathbf{y}) > 0$;
- **Aperiodic**: for every $\mathbf{x}$, $\gcd \{t \geqslant 1 \;:\; P^t(x, x) > 0\} = 1$;

Markov chains

...in one slide

$\{X_t \,:\, t \in \mathbb{N}\}$, (Ω, P).

- Irreducible: for every $\mathbf{x}, \mathbf{y} \in \Omega$, $\exists t \in \mathbb{N} \;:\; P^t(\mathbf{x}, \mathbf{y}) > 0$;
- Aperiodic: for every $\mathbf{x}$, gcd $\{t \geqslant 1 \,:\, P^t(x, x) > 0\} = 1$;
- **Stationary distribution**: $\pi \in \Delta(\Omega)$, $\pi P = \pi$.

Markov chains

...in one slide

$\{X_t \,:\, t \in \mathbb{N}\}$, (Ω, P).

- Irreducible: for every $\mathbf{x}, \mathbf{y} \in \Omega$, $\exists t \in \mathbb{N} \;:\; P^t(\mathbf{x}, \mathbf{y}) > 0$;
- Aperiodic: for every $\mathbf{x}$, $\gcd\{t \geqslant 1 \,:\, P^t(x, x) > 0\} = 1$;
- Stationary distribution: $\pi \in \Delta(\Omega)$, $\pi P = \pi$.

Irreducible + Aperiodic = Ergodic $\Longrightarrow$

$\Longrightarrow \pi$ is unique and $P^t(\mathbf{x}, \cdot) \to \pi$

Markov chains

...in one slide

$\{X_t \,:\, t \in \mathbb{N}\}$, (Ω, P).

- Irreducible: for every $\mathbf{x}, \mathbf{y} \in \Omega$, $\exists t \in \mathbb{N} \,:\, P^t(\mathbf{x}, \mathbf{y}) > 0$;
- Aperiodic: for every $\mathbf{x}$, $\gcd\{t \geqslant 1 \,:\, P^t(x, x) > 0\} = 1$;
- Stationary distribution: $\pi \in \Delta(\Omega)$, $\pi P = \pi$.

Irreducible + Aperiodic = Ergodic $\Longrightarrow$

$\Longrightarrow \pi$ is unique and $P^t(\mathbf{x}, \cdot) \to \pi$

- **Total variation distance** $\mu, \nu \in \Delta(\Omega)$

$$\|\mu - \nu\| = \max_{A \subseteq \Omega} |\mu(A) - \nu(A)| = \frac{1}{2} \sum_{\mathbf{x} \in \Omega} |\mu(\mathbf{x}) - \nu(\mathbf{x})|$$

Markov chains

...in one slide

$\{X_t \, : \, t \in \mathbb{N}\}$, (Ω, P).

- Irreducible: for every $\mathbf{x}, \mathbf{y} \in \Omega$, $\exists t \in \mathbb{N} \, : \, P^t(\mathbf{x}, \mathbf{y}) > 0$;
- Aperiodic: for every $\mathbf{x}$, $\gcd\{t \geqslant 1 \, : \, P^t(x, x) > 0\} = 1$;
- Stationary distribution: $\pi \in \Delta(\Omega)$, $\pi P = \pi$.

Irreducible + Aperiodic = Ergodic $\Longrightarrow$

$\Longrightarrow \pi$ is unique and $P^t(\mathbf{x}, \cdot) \to \pi$

- Total variation distance $\mu, \nu \in \Delta(\Omega)$

$$\|\mu - \nu\| = \max_{A \subseteq \Omega} |\mu(A) - \nu(A)| = \frac{1}{2} \sum_{\mathbf{x} \in \Omega} |\mu(\mathbf{x}) - \nu(\mathbf{x})|$$

- **Mixing Time**

$$t_{\text{mix}}(\varepsilon) = \min\{t \in \mathbb{N} \, : \, \|P^t(\mathbf{x}, \cdot) - \pi\| \leqslant \varepsilon \text{ for all } \mathbf{x} \in \Omega\}$$

Logit Dynamics

Definition

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. Logit dynamics for G is the Markov chain with state space $\Omega = S_1 \times \cdots \times S_n$ and transition matrix

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^{n} \frac{e^{\beta u_i(\mathbf{x}_{-i}, y_i)}}{T_i(\mathbf{x})} \mathbb{I}_{\{y_j = x_j \text{ for every } j \neq i\}}$$

Logit Dynamics

Definition

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. Logit dynamics for G is the Markov chain with state space $\Omega = S_1 \times \cdots \times S_n$ and transition matrix

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^{n} \frac{e^{\beta u_i(\mathbf{x}_{-i}, y_i)}}{T_i(\mathbf{x})} \mathbb{I}_{\{y_j = x_j \text{ for every } j \neq i\}}$$

Logit Dynamics defines an ergodic Markov chain

- What is the **stationary distribution** π?

Logit Dynamics

Definition

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. Logit dynamics for G is the Markov chain with state space $\Omega = S_1 \times \cdots \times S_n$ and transition matrix

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^{n} \frac{e^{\beta u_i(\mathbf{x}_{-i}, y_i)}}{T_i(\mathbf{x})} \mathbb{I}_{\{y_j = x_j \text{ for every } j \neq i\}}$$

Logit Dynamics defines an ergodic Markov chain

- What is the stationary distribution π?
- What is the **stationary expected social welfare** $\mathbf{E}_\pi [W]$?

Logit Dynamics

Definition

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. Logit dynamics for G is the Markov chain with state space $\Omega = S_1 \times \cdots \times S_n$ and transition matrix

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^{n} \frac{e^{\beta u_i(\mathbf{x}_{-i}, y_i)}}{T_i(\mathbf{x})} \mathbb{I}_{\{y_j = x_j \text{ for every } j \neq i\}}$$

Logit Dynamics defines an ergodic Markov chain

- What is the stationary distribution π?
- What is the stationary expected social welfare $\mathbf{E}_\pi [W]$?
- **How long it takes to get close to the stationary distribution?**

Logit Dynamics

Some recent related works

[Montanari, Saberi, FOCS'09]:
Hitting time of the best Nash equilibrium in a classical "game" (Ising model for ferromagnetism. Applications to the spread of innovations in a network)

Logit Dynamics

Some recent related works

[Montanari, Saberi, FOCS'09]:
Hitting time of the best Nash equilibrium in a classical "game" (Ising model for ferromagnetism. Applications to the spread of innovations in a network)

[Asadpour, Saberi, WINE'09]:
Hitting time of the *neighborhood* of best Nash equilibria for Atomic Selfish Routing and Load Balancing.

Logit Dynamics

Some recent related works

[Montanari, Saberi, FOCS'09]:
Hitting time of the best Nash equilibrium in a classical "game" (Ising model for ferromagnetism. Applications to the spread of innovations in a network)

[Asadpour, Saberi, WINE'09]:
Hitting time of the *neighborhood* of best Nash equilibria for Atomic Selfish Routing and Load Balancing.

They study **hitting time of Nash equilibria**. We propose **stationary distribution as equilibrium concept**.

Logit Dynamics

Some recent related works

[Montanari, Saberi, FOCS'09]:
Hitting time of the best Nash equilibrium in a classical "game" (Ising model for ferromagnetism. Applications to the spread of innovations in a network)

[Asadpour, Saberi, WINE'09]:
Hitting time of the *neighborhood* of best Nash equilibria for Atomic Selfish Routing and Load Balancing.

They study hitting time of Nash equilibria. We propose stationary distribution as equilibrium concept.

Nash equilibria	**Stationary distribution** of logit dynamics
Not Unique	**Unique**

Logit Dynamics

Some recent related works

[Montanari, Saberi, FOCS'09]:
Hitting time of the best Nash equilibrium in a classical "game" (Ising model for ferromagnetism. Applications to the spread of innovations in a network)

[Asadpour, Saberi, WINE'09]:
Hitting time of the *neighborhood* of best Nash equilibria for Atomic Selfish Routing and Load Balancing.

They study hitting time of Nash equilibria. We propose stationary distribution as equilibrium concept.

Nash equilibria	**Stationary distribution** of logit dynamics
Not Unique	Unique
Local	**Global**

Example: Matching Pennies

- If both choose the same side then Player 1 wins;
- If they choose different sides then Player 2 wins

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

Example: Matching Pennies

- If both choose the same side then Player 1 wins;
- If they choose different sides then Player 2 wins

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

No pure Nash, one mixed Nash $\sigma_1 = \sigma_2 = (1/2, 1/2)$

Example: Matching Pennies

- If both choose the same side then Player 1 wins;
- If they choose different sides then Player 2 wins

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

No pure Nash, one mixed Nash $\sigma_1 = \sigma_2 = (1/2, 1/2)$

Logit dynamics

Example:

$$\sigma_1(H \,|\, (-, T)) \quad =$$

Example: Matching Pennies

- If both choose the same side then Player 1 wins;
- If they choose different sides then Player 2 wins

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

No pure Nash, one mixed Nash $\sigma_1 = \sigma_2 = (1/2, 1/2)$

Logit dynamics

Example:

$$\sigma_1(H \mid (-, T)) = \frac{e^{\beta u_1(H,T)}}{e^{\beta u_1(H,T)} + e^{\beta u_1(T,T)}}$$

Example: Matching Pennies

- If both choose the same side then Player 1 wins;
- If they choose different sides then Player 2 wins

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

No pure Nash, one mixed Nash $\sigma_1 = \sigma_2 = (1/2, 1/2)$

Logit dynamics

Example:

$$\begin{aligned} \sigma_1(H \mid (-, T)) &= \frac{e^{\beta u_1(H,T)}}{e^{\beta u_1(H,T)} + e^{\beta u_1(T,T)}} \\ &= \frac{e^{-\beta}}{e^{-\beta} + e^{\beta}} \end{aligned}$$

Example: Matching Pennies

- If both choose the same side then Player 1 wins;
- If they choose different sides then Player 2 wins

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

No pure Nash, one mixed Nash $\sigma_1 = \sigma_2 = (1/2, 1/2)$

Logit dynamics

Example:

$$\begin{aligned} \sigma_1(H \mid (-, T)) &= \frac{e^{\beta u_1(H,T)}}{e^{\beta u_1(H,T)} + e^{\beta u_1(T,T)}} \\ &= \frac{e^{-\beta}}{e^{-\beta} + e^{\beta}} = \frac{1}{1 + e^{2\beta}} \end{aligned}$$

Matching Pennies

Logit dynamics

$$P = \frac{1}{2}\left(\begin{array}{c|cccc} & HH & HT & TH & TT \\ \hline HH & 1 & b & (1-b) & 0 \\ HT & (1-b) & 1 & 0 & b \\ TH & b & 0 & 1 & (1-b) \\ TT & 0 & (1-b) & b & 1 \end{array}\right)$$

$b = \frac{1}{1+e^{-2\beta}}$.

Matching Pennies

Logit dynamics

$$P = \frac{1}{2} \left(\begin{array}{c|cccc} & HH & \mathbf{HT} & TH & TT \\ \hline HH & 1 & b & (1-b) & 0 \\ HT & (1-b) & 1 & 0 & b \\ TH & b & 0 & 1 & (1-b) \\ \mathbf{TT} & 0 & \mathbf{(1\text{-}b)} & b & 1 \end{array} \right)$$

$b = \frac{1}{1+e^{-2\beta}}$.

Matching Pennies
Logit dynamics

$$P = \frac{1}{2} \begin{array}{c|cccc} & HH & HT & TH & TT \\ \hline HH & 1 & b & (1-b) & 0 \\ HT & (1-b) & 1 & 0 & b \\ TH & b & 0 & 1 & (1-b) \\ TT & 0 & (1-b) & b & 1 \end{array}$$

$b = \frac{1}{1+e^{-2\beta}}$.

- $\pi = \frac{1}{4}(1, 1, 1, 1)$

$$\mathbf{E}_\pi [W] = 0$$

Matching Pennies

Logit dynamics

$$P = \frac{1}{2} \left(\begin{array}{c|cccc} & HH & HT & TH & TT \\ \hline HH & 1 & b & (1-b) & 0 \\ HT & (1-b) & 1 & 0 & b \\ TH & b & 0 & 1 & (1-b) \\ TT & 0 & (1-b) & b & 1 \end{array} \right)$$

$b = \frac{1}{1+e^{-2\beta}}$.

- $\pi = \frac{1}{4}(1, 1, 1, 1)$

$$\mathbf{E}_\pi [W] = 0$$

- $\|P^3(\mathbf{x}, \cdot) - \pi\| < 1/2$ so

$$t_{\text{mix}} = \mathcal{O}(1)$$

(upper bounded by a constant for every β)

Logit Dynamics

Potential games

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. $\Phi : S_1 \times \cdots \times S_n \to \mathbb{R}$ **exact potential** if for every profile $\mathbf{x}$, for every player i, and for every action y

$$u_i(\mathbf{x}_{-i}, y) - u_i(\mathbf{x}) = -\left[\Phi(\mathbf{x}_{-i}, y) - \Phi(\mathbf{x})\right]$$

Logit Dynamics

Potential games

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. $\Phi : S_1 \times \cdots \times S_n \to \mathbb{R}$ **exact potential** if for every profile $\mathbf{x}$, for every player i, and for every action y

$$u_i(\mathbf{x}_{-i}, y) - u_i(\mathbf{x}) = -\left[\Phi(\mathbf{x}_{-i}, y) - \Phi(\mathbf{x})\right]$$

Logit dynamics for potential games

The stationary distribution is the Gibbs one

$$\pi(\mathbf{x}) = \frac{e^{-\beta\Phi(\mathbf{x})}}{Z}$$

Logit Dynamics

Potential games

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. $\Phi : S_1 \times \cdots \times S_n \to \mathbb{R}$ **exact potential** if for every profile $\mathbf{x}$, for every player i, and for every action y

$$u_i(\mathbf{x}_{-i}, y) - u_i(\mathbf{x}) = -\left[\Phi(\mathbf{x}_{-i}, y) - \Phi(\mathbf{x})\right]$$

Logit dynamics for potential games

The stationary distribution is the Gibbs one

$$\pi(\mathbf{x}) = \frac{e^{-\beta\Phi(\mathbf{x})}}{Z}$$

Observation [Blume'93]

For $\beta \to \infty$ the stationary distribution π is concentrated over the global minima of the potential function.

Logit Dynamics

Potential games

$\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$. $\Phi : S_1 \times \cdots \times S_n \to \mathbb{R}$ **exact potential** if for every profile $\mathbf{x}$, for every player i, and for every action y

$$u_i(\mathbf{x}_{-i}, y) - u_i(\mathbf{x}) = -\left[\Phi(\mathbf{x}_{-i}, y) - \Phi(\mathbf{x})\right]$$

Logit dynamics for potential games

The stationary distribution is the Gibbs one

$$\pi(\mathbf{x}) = \frac{e^{-\beta\Phi(\mathbf{x})}}{Z}$$

Observation [Blume'93]

For $\beta \to \infty$ the stationary distribution π is concentrated over the global minima of the potential function.

$$\Longrightarrow \frac{\mathrm{opt}(W)}{\mathbf{E}_\pi[W]} \to \text{ Price of Stability}$$

Chicken Game

Description

- If both STOP then none wins;
- If both PASS then both lose;
- If one PASSes and one STOP, then who passes win.

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

Chicken Game

Description

- If both STOP then none wins;
- If both PASS then both lose;
- If one PASSes and one STOP, then who passes win.

	S	P
S	$(0,0)$	$\mathbf{(0,\ 1)}$
P	$\mathbf{(1,\ 0)}$	$(-1,-1)$

Two pure Nash, one mixed Nash.
Potential game: stationary distribution is for free.

Chicken Game

Description

- If both STOP then none wins;
- If both PASS then both lose;
- If one PASSes and one STOP, then who passes win.

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

Two pure Nash, one mixed Nash.
Potential game: stationary distribution is for free.

Stationary expected social welfare

$$\pi(SS) = \pi(PP) = \frac{1}{2(1+e^{\beta})}$$
$$\pi(SP) = \pi(PS) = \frac{1}{2(1+e^{-\beta})}$$

$$\mathbf{E}_{\pi}[W] = \frac{e^{\beta}-1}{e^{\beta}+1}$$

Chicken Game

Description

- If both STOP then none wins;
- If both PASS then both lose;
- If one PASSes and one STOP, then who passes win.

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

Two pure Nash, one mixed Nash.
Potential game: stationary distribution is for free.

Stationary expected social welfare

$$\pi(SS) = \pi(PP) = \frac{1}{2(1+e^{\beta})}$$
$$\pi(SP) = \pi(PS) = \frac{1}{2(1+e^{-\beta})}$$

$$\mathbf{E}_{\pi}[W] = \frac{e^{\beta}-1}{e^{\beta}+1}$$

Observations

- Expected social welfare **tends to** 1 for $\beta \to \infty$;

Chicken Game

Description

- If both STOP then none wins;
- If both PASS then both lose;
- If one PASSes and one STOP, then who passes win.

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

Two pure Nash, one mixed Nash.
Potential game: stationary distribution is for free.

Stationary expected social welfare

$$\pi(SS) = \pi(PP) = \frac{1}{2(1+e^{\beta})}$$
$$\pi(SP) = \pi(PS) = \frac{1}{2(1+e^{-\beta})}$$

$$\mathbf{E}_{\pi}[W] = \frac{e^{\beta}-1}{e^{\beta}+1}$$

Observations

- Expected social welfare tends to 1 for $\beta \to \infty$;
- Expected social welfare is **fair**.

Chicken Game

Mixing time

Mixing time is **exponential** in β

$$t_{\text{mix}} = \Theta(e^{\beta})$$

Chicken Game

Mixing time

Mixing time is exponential in β

$$t_{\text{mix}} = \Theta(e^{\beta})$$

Intuition

- The two Nash equilibria have the same stationary probability;
- When β is *large*, it takes a long time to go from one Nash equilibrium to the other one.

Chicken Game

Mixing time

Mixing time is exponential in β

$$t_{\text{mix}} = \Theta(e^{\beta})$$

Intuition

- The two Nash equilibria have the same stationary probability;
- When β is *large*, it takes a long time to go from one Nash equilibrium to the other one.

Generalization

Chicken Game is an anti-coordination game. The results extend to all **coordination** and **anti-coordination** games.

OR-game

Definition

A **trivial** game with a **non-trivial** mixing time analysis.

OR-game

Definition

A **trivial** game with a **non-trivial** mixing time analysis.

Every player has two strategies, say $\{0, 1\}$, and each player pays the OR of the strategies of all players (including herself).

$$u_i(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} = \mathbf{0}; \\ -1 & \text{otherwise.} \end{cases}$$

OR-game

Definition

A **trivial** game with a **non-trivial** mixing time analysis.

Every player has two strategies, say $\{0,1\}$, and each player pays the OR of the strategies of all players (including herself).

$$u_i(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} = \mathbf{0}; \\ -1 & \text{otherwise.} \end{cases}$$

Almost every profile is Nash equilibrium (except profiles with exactly one 1)

$$W(\text{best Nash eq.}) = 0; \qquad W(\text{worst Nash eq.}) = -n$$

OR-game

Definition

A **trivial** game with a **non-trivial** mixing time analysis.

Every player has two strategies, say $\{0, 1\}$, and each player pays the OR of the strategies of all players (including herself).

$$u_i(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} = \mathbf{0}; \\ -1 & \text{otherwise.} \end{cases}$$

Almost every profile is Nash equilibrium (except profiles with exactly one 1)

$$W(\text{best Nash eq.}) = 0; \qquad W(\text{worst Nash eq.}) = -n$$

Expected social welfare

$$\mathbf{E}_{\pi}[W] = -\frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}} \cdot n$$

OR-game

Definition

A **trivial** game with a **non-trivial** mixing time analysis.

Every player has two strategies, say $\{0, 1\}$, and each player pays the OR of the strategies of all players (including herself).

$$u_i(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} = \mathbf{0}; \\ -1 & \text{otherwise.} \end{cases}$$

Almost every profile is Nash equilibrium (except profiles with exactly one 1)

$$W(\text{best Nash eq.}) = 0; \qquad W(\text{worst Nash eq.}) = -n$$

Expected social welfare

$$\mathbf{E}_\pi[W] = -\frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}} \cdot n \qquad \mathbf{E}_\pi[W] \to 0 \quad \text{for } \beta \to \infty$$

OR-game

Mixing time

$$t_{\text{mix}} \approx \begin{cases} \mathcal{O}(n \log n) & \text{if } \beta < \log n \\ & \\ & \end{cases}$$

OR-game

Mixing time

$$t_{\text{mix}} \approx \begin{cases} \mathcal{O}(n \log n) & \text{if } \beta < \log n \\ \Theta(n^c) & \text{if } \beta = c \log n, c > 1 \text{ constant} \\ \end{cases}$$

OR-game

Mixing time

$$t_{\text{mix}} \approx \begin{cases} \mathcal{O}(n \log n) & \text{if } \beta < \log n \\ \Theta(n^c) & \text{if } \beta = c \log n, c > 1 \text{ constant} \\ \Theta(2^n) & \text{otherwise} \end{cases}$$

OR-game

Mixing time

$$t_{\text{mix}} \approx \begin{cases} \mathcal{O}(n \log n) & \text{if } \beta < \log n \\ \Theta(n^c) & \text{if } \beta = c \log n, c > 1 \text{ constant} \\ \Theta(2^n) & \text{otherwise} \end{cases}$$

Intuition

- When β is *large* the stationary distribution is concentrated in profile $\mathbf{0}$;
- From a profile with at least two 1's, every players choose u.a.r. over $\{0, 1\}$

OR-game

Mixing time

$$t_{\text{mix}} \approx \begin{cases} \mathcal{O}(n \log n) & \text{if } \beta < \log n \\ \Theta(n^c) & \text{if } \beta = c \log n, c > 1 \text{ constant} \\ \Theta(2^n) & \text{otherwise} \end{cases}$$

Intuition

- When β is *large* the stationary distribution is concentrated in profile $\mathbf{0}$;
- From a profile with at least two 1's, every players choose u.a.r. over $\{0, 1\}$

Proof techniques:

- Path-coupling for the upper bound;
- Bottleneck ratio for the lower bound.

Further and Future Investigations

- **Logit dynamics for more interesting (class of) games**;
  - Potential games: Mixing time depends on the *shape* of the potential function
    (Lipschitz conditions, Number of local minima, Maximum *Slope*, ...);
  - Stationary expected social welfare vs PoA / PoS for classical games;

Further and Future Investigations

- Logit dynamics for more *interesting* (class of) games;
  - Potential games: Mixing time depends on the *shape* of the potential function
    (Lipschitz conditions, Number of local minima, Maximum *Slope*, ...);
  - Stationary expected social welfare vs PoA / PoS for classical games;
- **Other randomized best response dynamics**
  - E.g. All players play simultaneously;

Further and Future Investigations

- Logit dynamics for more *interesting* (class of) games;
  - Potential games: Mixing time depends on the *shape* of the potential function
    (Lipschitz conditions, Number of local minima, Maximum *Slope*, . . .);
  - Stationary expected social welfare vs PoA / PoS for classical games;
- Other randomized best response dynamics
  - E.g. All players play simultaneously;
- **Connections with other disciplines**
  - Statistical Physics;
  - Evolutionary biology??

Further and Future Investigations

- Logit dynamics for more *interesting* (class of) games;
  - Potential games: Mixing time depends on the *shape* of the potential function (Lipschitz conditions, Number of local minima, Maximum *Slope*, ...);
  - Stationary expected social welfare vs PoA / PoS for classical games;
- Other randomized best response dynamics
  - E.g. All players play simultaneously;
- Connections with other disciplines
  - Statistical Physics;
  - Evolutionary biology??
- **What happens when mixing time is exponential?**
  - Chaotic behavior during transient phase?
  - Polynomially *almost* stationary distributions?

References

Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano.
Mixing time and Stationary Expected Social Welfare of Logit Dynamics.
Submitted, 2010 (http://arxiv.org/abs/1002.3474).

References

Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano.
Mixing time and Stationary Expected Social Welfare of Logit Dynamics.
Submitted, 2010 (http://arxiv.org/abs/1002.3474).

Thank you!