Logit dynamics for strategic games

Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano
Dipartimento di Informatica "Renato M. Capocelli", Università di Salerno, ITALY

We study a "noisy" best-response dynamics for strategic games. We propose the stationary distribution of the induced Markov chain as equilibrium concept for systems where agents have limited rationality and analyze the convergence time.

Motivations

Agents often do not behave rationally.

Possible reasons
- Limited computational power;
- Incomplete information;
- Intrinsic irrationality.

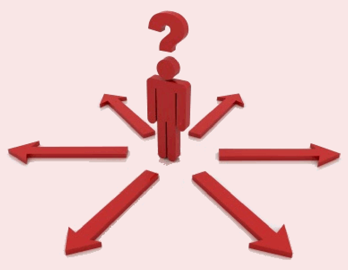

Standard game-theoretic predictions might be unsatisfactory.

Best response dynamics and Nash equilibria
Not unique; Hard to compute; Long time to reach.

Model

Logit dynamics [Blume, Games and Economic Behavior, 1993]
$\mathcal{G} = ([n], (S_1, \ldots, S_n), (u_1, \ldots, u_n))$ strategic game.
When at profile $\mathbf{x} = (x_1, \ldots, x_n)$, pick a player $i \in [n]$ u.a.r. and update her strategy at $y \in S_i$ with probability proportional to $\exp \beta u_i(\mathbf{x}_{-i}, y)$

Logit dynamics models situations where players are only *approximately* rational.

Rationality level
Parameter β is the *inverse noise* (or *inverse temperature*):
- $\beta = 0$: players choose u.a.r. (no rationality);
- $\beta > 0$: bias toward strategies promising higher payoffs;
- $\beta \to \infty$: players best-respond (full rationality).

For logit dynamics equilibrium exists and is unique...

Existence and uniqueness
Logit dynamics for $\mathcal{G}$ defines a family of ergodic Markov chains over the set of strategy profiles $S_1 \times \cdots \times S_n$

...and for some classes of games it is given by an explicit formula.

Potential games
If $\mathcal{G}$ is a potential game with potential function Φ:
- Logit dynamics is *reversible*;
- Stationary distribution = Gibbs:

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

- Logit dynamics for $\mathcal{G} \equiv$ Glauber dynamics for π.

To analyze logit dynamics we need tools that have been used mainly in statistical mechanics.

Goals

We want to measure the performance of the system at equilibrium and ...

Stationary expected social welfare
Expected social welfare when profiles are random according to the stationary distribution

$$\mathbf{E}_\pi[W] = \Sigma_{\mathbf{x} \in S_1 \times \cdots \times S_n} W(\mathbf{x})\pi(\mathbf{x})$$

...how long it takes to reach it.

Mixing time
The time it takes to get close to the stationary distribution

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

in total variation distance

$$\|P^t(\mathbf{x}, \cdot) - \pi\| = \frac{1}{2}\Sigma_{\mathbf{y} \in S_1 \times \cdots \times S_n}|P^t(\mathbf{x}, \mathbf{y}) - \pi(\mathbf{y})|$$

Examples

***Matching pennies*: convergence is fast, and stationary social welfare agrees with social welfare at (mixed) Nash.**

Matching pennies

	Head	Tail
Head	$(1, -1)$	$(-1, 1)$
Tail	$(-1, 1)$	$(1, -1)$

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

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

- Stationary distribution: $\pi = \frac{1}{4}(1, 1, 1, 1)$
- Stationary expected social welfare: $\mathbf{E}_\pi[W] = 0$
- Mixing time: $t_{\text{mix}} = 3$

***Chicken game*: convergence time increases with level of rationality. Stationary social welfare is smaller than social welfare at Nash equilibria, but it is *fair* (equal for both players).**

Chicken game

	Stop	Pass
Stop	$(0, 0)$	$(0, 1)$
Pass	$(1, 0)$	$(-1, -1)$

$$\frac{1}{2}\begin{array}{c|cccc} & SS & SP & PS & PP \\ \hline SS & 2b & (1-b) & (1-b) & 0 \\ SP & b & 2(1-b) & 0 & b \\ PS & b & 0 & 2(1-b) & b \\ PP & 0 & (1-b) & (1-b) & 2b \end{array}$$

- $\pi = \frac{1}{2}(b, 1-b, 1-b, b)$
- $\mathbf{E}_\pi[W] = \mathbf{E}_\pi[u_1] + \mathbf{E}_\pi[u_2] = (1 - 2b)$
- $t_{\text{mix}} = \Theta\left(e^{\beta}\right)$

Some results

There are games where the convergence time depends on the level of rationality...

Potential games

$$t_{\text{mix}} = \mathcal{O}\left(e^{\beta(\Phi_{\max} - \Phi_{\min})}\text{poly}(n)\right)$$

...games where the level of rationality does not affect it ...

Games with dominant strategies

$$t_{\text{mix}} = \mathcal{O}\left(f\left(\mathcal{G}\right)\right) \quad \text{regardless of } \beta$$

...and games where it depends on the *topology* of interactions

Graphical coordination games
- Clique: t_{mix} = exponential in β and n^2;
- Ring: t_{mix} = exponential in β.

Future research

- **Other games**: Study mixing time and stationary expected social welfare for other "important" classes of games;
- **Metastability**: Investigate what happens during the *transient* phase of the Markov chain, when the mixing time is large;
- **Other noisy dynamics**: Consider simultaneous updates, non-uniform noise, *structured* strategy sets;
- **Connections**: Explore further connections with other disciplines: physics, economics, biology, social sciences, ...

Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano
Mixing Time and Stationary Expected Social Welfare of Logit Dynamics
In *Proc. of 3rd SAGT*, LNCS 6386, pp. 54-65. Springer, 2010.

Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano
Convergence to Equilibrium of Logit Dynamics for Strategic Games
In *Proc. of 23rd SPAA*, to appear, 2011.