Consistency of Preferences and Near-Potential Games

Ozan Candogan, Ishai Menache,
Asu Ozdaglar and Pablo A. Parrilo

Laboratory for Information and Decision Systems
Electrical Engineering and Computer Science
Massachusetts Institute of Technology

LCCC - Lund - March 2010

Outline

Motivation

- Potential games are games in which preferences of all players are aligned with a global objective.
  - easy to analyze
  - pure Nash equilibrium exists
  - simple dynamics converge to an equilibrium
- How "close" is a game to a potential game?
- What is the topology of the space of preferences?
- Are there "natural" decompositions of games?
- How to modify a game to make it potential?
- Useful as analysis methodology, but also for game design.

Main Contributions

- Analysis of the global structure of preferences
- Canonical decomposition: potential and harmonic components
- Projection schemes to find the components.
- Closed form solutions to the projection problem.
- Equilibria of a game are ϵ-equilibria of its projection, and equilibria of the projected game are ϵ equilibria of the initial game.
- Analysis of games in terms of their components

Potential Games

- We consider finite games in strategic form, $\mathcal{G} = \langle \mathcal{M}, \{E^m\}_{m \in \mathcal{M}}, \{u^m\}_{m \in \mathcal{M}} \rangle$.
- $\mathcal{G}$ is an exact potential game if $\exists\, \Phi : E \to \mathbb{R}$ such that

$$\Phi(x^m, x^{-m}) - \Phi(y^m, x^{-m}) = u^m(x^m, x^{-m}) - u^m(y^m, x^{-m}),$$

- Weaker notion: ordinal potential game, if the utility differences above agree only in sign.
- Potential Φ aggregates and explains incentives of all players.
- Examples: congestion games, etc.

Potential Games

- A global maximum of an ordinal potential is a pure Nash equilibrium.
- Every finite potential game has a pure equilibrium.
- Many learning dynamics (such as better-reply dynamics, fictitious play, spatial adaptive play) “converge” to a pure Nash equilibrium [Monderer and Shapley 96], [Young 98], [Marden, Arslan, Shamma 06, 07].

Potential Games

- When is a given game a potential game?
- More important, what are the obstructions, and what is the underlying structure?

Existence of Exact Potential

A path is a collection of strategy profiles $\gamma = (x_0, \ldots, x_N)$ such that x_i and x_{i+1} differ in the strategy of exactly one player where $x_i \in E$ for $i \in \{0, 1, \ldots N\}$. For any path γ, let

$$I(\gamma) = \sum_{i=1}^{N} u^{m_i}(x_i) - u^{m_i}(x_{i-1}),$$

where m_i denotes the player changing its strategy in the ith step of the path.

Theorem ([Monderer and Shapley 96])

A game $\mathcal{G}$ is an exact potential game if and only if for all simple closed paths, γ, $I(\gamma) = 0$. Moreover, it is sufficient to check closed paths of length 4.

Existence of Exact Potential

- Let $I(\gamma) \neq 0$, if potential existed then it would increase when the cycle is completed.
- The condition for existence of exact potential is linear. The set of exact potential games is a subspace of the space of games.
- The set of exact potential games is "small".

Theorem

Consider games with set of players $\mathcal{M}$, and joint strategy space $E = \prod_{m \in \mathcal{M}} E^m$

1. *The dimension of the space of games is $|\mathcal{M}| \prod_{m \in \mathcal{M}} |E^m|$.*
2. *The dimension of the subspace of exact potential games is*

$$\prod_{m \in \mathcal{M}} |E^m| + \sum_{m \in \mathcal{M}} \prod_{k \in \mathcal{M}, k \neq m} |E^k| - 1.$$

Existence of Ordinal Potential

- A weak improvement cycle is a cycle for which at each step of which the utility of the player whose strategy is modified is nondecreasing (and at least at one step the change is strictly positive).
- A game is an ordinal potential game if and only if it contains no weak improvement cycles [Voorneveld and Norde 97].

Game Flows: 3-Player Example

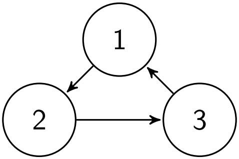

- $E^m = \{0, 1\}$ for all $m \in \mathcal{M}$, and payoff of player i be -1 if its strategy is the same with its successor, 0 otherwise.
- This game is neither an exact nor an ordinal potential game.

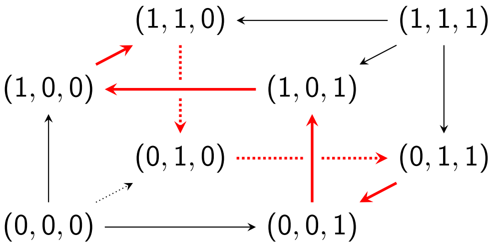

Global Structure of Preferences

- What is the global structure of these cycles?
- Equivalently, topological structure of aggregated preferences.
- Conceptually similar to structure of (continuous) vector fields.
- A well-developed theory from algebraic topology, we need the combinatorial analogue.

Helmholtz (Hodge) Decomposition

The Helmholtz Decomposition allows orthogonal decomposition of a vector field into three vector fields:

- Gradient flow (globally acyclic component)
- Harmonic flow (locally acyclic but globally cyclic component)
- Curl flow (locally cyclic component).

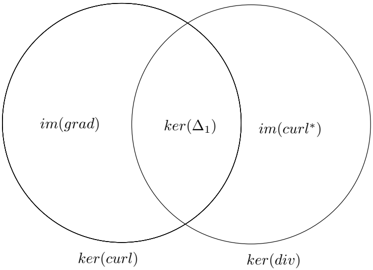

Figure: Helmholtz Decomposition

Helmholtz decomposition (a cartoon)

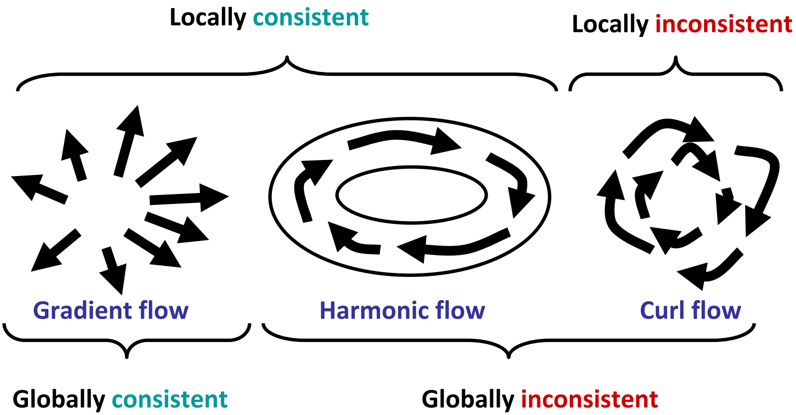

Redefining Potential Games

- For all $m \in \mathcal{M}$ let $W^m : E \times E \to \mathbb{R}$ satisfy

$$W^m(\mathbf{p}, \mathbf{q}) = \begin{cases} 1 & \text{if } \mathbf{p}, \mathbf{q} \text{ differ in the strategy of player } m \text{ only} \\ 0 & \text{otherwise.} \end{cases}$$

- For all $m \in \mathcal{M}$, define a difference operator D_m such that,

$$(D_m\phi)(\mathbf{p}, \mathbf{q}) = W^m(\mathbf{p}, \mathbf{q}) \left(\phi(\mathbf{q}) - \phi(\mathbf{p})\right).$$

  where $\mathbf{p}, \mathbf{q} \in E$ and $\phi : E \to \mathbb{R}$.

- Note that a game is an exact potential game if and only if

$$D_m u^m = D_m \phi$$

  for all $m \in \mathcal{M}$.

Redefining Potential Games

- $\delta_0 = \sum_{m \in \mathcal{M}} D_m$ is a combinatorial gradient operator.
- Image spaces of operators D_m $m \in \mathcal{M}$ are orthogonal
- A game is an exact potential game if and only if

$$\sum_{m \in \mathcal{M}} D_m u^m = \sum_{m \in \mathcal{M}} D_m \phi = \delta_0 \phi.$$

Exact Potential Games - Alternative Definition

A game is an exact potential game if and only if $\sum_{m \in \mathcal{M}} D_m u^m$ is a gradient flow.

Decomposition: Potential, Harmonic, and Nonstrategic

Decomposition of game flows induces a similar partition of the space of games:

- When going from utilities to flows, the nonstrategic component is removed.
- If we start from utilities (not preferences), always locally consistent.
- Therefore, two flow components: potential and harmonic

Thus, the space of games has a canonical direct sum decomposition:

$$G = G_{\text{potential}} \oplus G_{\text{harmonic}} \oplus G_{\text{nonstrategic}},$$

where the components are orthogonal subspaces.

Bimatrix games

For two-player games, simple explicit formulas.
Assume the game is given by matrices (A, B), and (for simplicity), the non-strategic component is zero (i.e., $\mathbf{1}^T A = 0, B\mathbf{1} = 0$). Define

$$S := \frac{1}{2}(A + B), \quad D := \frac{1}{2}(A - B), \quad \Gamma := \frac{1}{2n}(A\mathbf{1}\mathbf{1}^T - \mathbf{1}\mathbf{1}^T B).$$

- Potential component:

$$(S + \Gamma, \quad S - \Gamma)$$

- Harmonic component:

$$(D - \Gamma, \quad -D + \Gamma)$$

Notice that the harmonic component is zero sum.

Harmonic games

Very different properties than potential games.
Agreement between players is never a posibility!

- Simple examples: rock-paper-scissors, cyclic games, etc.
- Essentially, sums of cycles.
- Generically, *never* have pure Nash equilibria.
- Uniformly mixed profile (for all players) is mixed Nash.

Other interesting static and dynamic properties (e.g., correlated equilibria, best-response dynamics, etc.)

Projection on the Set of Exact Potential Games

- We solve,

$$d^2(\mathcal{G}) = \min_{\phi \in C_0} \quad ||\delta_0 \phi - \sum_{m \in \mathcal{M}} D_m u^m||_2^2,$$

  to find a potential function that best represents a given collection of utilities (C_0 is the space of real valued functions defined on E) .

- The utilities that represent the potential and that are close to initial utilities can be constructed by solving an additional optimization problem (for a fixed ϕ, and for all $m \in \mathcal{M}$):

$$\begin{aligned} \hat{u}^m = \arg\min_{\bar{u}^m} \quad & ||u^m - \bar{u}^m||_2^2 \\ s.t. \quad & D_m \bar{u}^m = D_m \phi \\ & \bar{u}^m \in C_0. \end{aligned}$$

Projection on the Set of Exact Potential Games

Theorem

If all players have same number of strategies, the optimal projection is given in closed form by

$$\phi = \left(\sum_{m \in \mathcal{M}} \Pi_m\right)^{\dagger} \sum_{m \in \mathcal{M}} \Pi_m u^m,$$

and

$$\hat{u}^m = (I - \Pi_m)u^m + \Pi_m \left(\sum_{k \in \mathcal{M}} \Pi_k\right)^{\dagger} \sum_{k \in \mathcal{M}} \Pi_k u_k.$$

Here $\Pi_m = D_m^* D_m$ is the projection operator to the orthogonal complement of the kernel of D_m ($*$ denotes the adjoint of an operator).

Projection on the Set of Exact Potential Games

- For any $m \in \mathcal{M}$, $\Pi_m u^m$ and $(I - \Pi_m) u^m$ are respectively the strategic and nonstrategic components of the utility of player m.
- ϕ solves,
$$\sum_{m \in \mathcal{M}} \Pi_m \phi = \sum_{m \in \mathcal{M}} \Pi_m u^m.$$
Hence, optimal ϕ is a function which represents the sum of strategic parts of utilities of different users.
- $\hat{u}^m$ is the sum of the nonstrategic part of u^m and the strategic part of the potential ϕ.

Consequences

Nice and beautiful. But (if that's not enough!) why should we care?

- Provides classes of games with simpler structures, for which stronger results can be proved.
- Yields a natural mechanism for approximation, for both static and dynamical properties.

Let's see this...

Equilibria of a Game and its Projection

Theorem

Let $\mathcal{G}$ be a game and $\hat{\mathcal{G}}$ be its projection. Any equilibrium of $\hat{\mathcal{G}}$ is an ϵ-equilibrium of $\mathcal{G}$ and any equilibrium of $\mathcal{G}$ is an ϵ-equilibrium of $\hat{\mathcal{G}}$ for $\epsilon \leq \sqrt{2} \cdot d(\mathcal{G})$.

- Provided that the projection distance is small, equilibria of the projected game are close to the equilibria of the initial game.

Simulation example

- Consider an average opinion game on a graph. Payoff of each player satisfies,

$$u^m(\mathbf{p}) = 2\hat{M} - (\hat{M}^m - \mathbf{p}^m)^2,$$

where $\hat{M}^m$ is the median of $\mathbf{p}^k$, $k \in N(m)$.

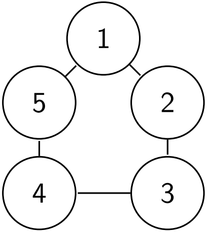

This game is not an exact (or ordinal) potential game.
With small perturbation in the payoffs, it can be projected to the set of potential games.

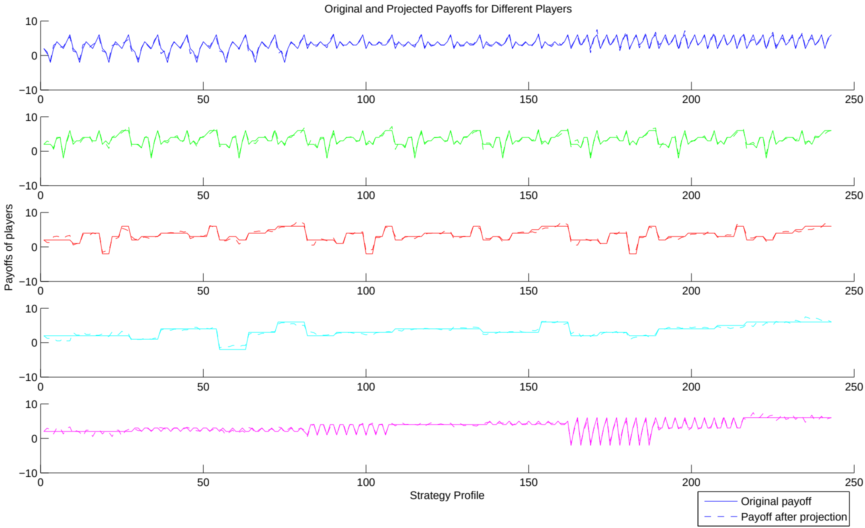

Original and Projected Payoffs for Different Players
Payoffs of players
Strategy Profile
10
0
−10
0
50
100
150
200
250
Original payoff
Payoff after projection

Wireless Power Control Application

- A set of mobiles (users) $\mathcal{M} = \{1, \ldots, M\}$ share the same wireless spectrum (single channel).
- We denote by $\mathbf{p} = (p_1, \ldots, p_M)$ the power allocation (vector) of the mobiles.
- Power constraints: $\mathcal{P}_m = \{p_m \mid \underline{P}_m \leq p_m \leq \bar{P}_m\}$, with $\underline{P}_m > 0$.
  - Upper bound represents a constraint on the maximum power usage
  - Lower bound represents a minimum QoS constraint for the mobile
- The rate (throughput) of user m is given by

$$r_m(\mathbf{p}) = \log\left(1 + \gamma \cdot \mathrm{SINR}_m(\mathbf{p})\right),$$

where, $\gamma > 0$ is the spreading gain of the CDMA system and

$$\mathrm{SINR}_m(\mathbf{p}) = \frac{h_{mm} p_m}{N_0 + \sum_{k \neq m} h_{km} p_k}.$$

Here, h_{km} is the channel gain between user k's transmitter and user m's receiver.

User Utilities and Equilibrium

- Each user is interested in maximizing a net rate-utility, which captures a tradeoff between the obtained rate and power cost:

$$u_m(\mathbf{p}) = r_m(\mathbf{p}) - \lambda_m p_m,$$

  where λ_m is a user-specific price per unit power.

- We refer to the induced game among the users as the power game and denote it by $\mathcal{G}$.
- Existence of a pure Nash equilibrium follows because the underlying game is a *concave game*.
- We are also interested in "approximate equilibria" of the power game, for which we use the concept of ϵ-(Nash) equilibria.
  - For a given ϵ, we denote by $\mathcal{I}_\epsilon$ the set of ϵ-equilibria of the power game $\mathcal{G}$, i.e.,

$$\mathcal{I}_\epsilon = \{\mathbf{p} \mid u_m(p_m, \mathbf{p}_{-m}) \geq u_m(q_m, \mathbf{p}_{-m}) - \epsilon, \quad \text{for all } m \in \mathcal{M}, \ q_m \in \mathcal{P}_m\}$$

System Utility

- Assume that a central planner wishes to impose a general performance objective over the network formulated as

$$\max_{\mathbf{p} \in \mathcal{P}} U_0(\mathbf{p}),$$

  where $\mathcal{P} = \mathcal{P}_1 \times \cdots \times \mathcal{P}_m$ is the joint feasible power set.
- We refer to $U_0(\cdot)$ as the system utility-function.
- We denote the optimal solution of this system optimization problem by p^* and refer to it as the desired operating point.
- Our goal is to set the prices such that the equilibrium of the power game can approximate the desired operating point p^*.

Potential Game Approximation

- We approximate the power game with a potential game.
- We consider a slightly modified game with player utility functions given by

$$\tilde{u}_m(\mathbf{p}) = \tilde{r}_m(\mathbf{p}) - \lambda_m p_m$$

  where $\tilde{r}_m(\mathbf{p}) = \log\left(\gamma \mathrm{SINR}_m(\mathbf{p})\right)$.
- We refer to this game as the potentialized game and denote it by $\tilde{\mathcal{G}} = \langle \mathcal{M}, \{\tilde{u}_m\}, \{\mathcal{P}_m\}\rangle$.
- For high-SINR regime (γ satisfies $\gamma \gg 1$ or $h_{mm} \gg h_{km}$ for all $k \neq m$),the modified rate formula $\tilde{r}_m(\mathbf{p}) \approx r_m(\mathbf{p})$ serves as a good approximation for the true rate, and thus $\tilde{u}_m(\mathbf{p}) \approx u_m(\mathbf{p})$.

Pricing in the Modified Game

Theorem

The modified game $\tilde{\mathcal{G}}$ is a potential game. The corresponding potential function is given by

$$\phi(\mathbf{p}) = \sum_m \log(p_m) - \lambda_m p_m.$$

- $\tilde{\mathcal{G}}$ has a unique equilibrium.
- The potential function suggests a simple linear pricing scheme.

Theorem

Let $\mathbf{p}^*$ be the desired operating point. Assume that the prices λ^* are given by

$$\lambda_m^* = \frac{1}{p_m^*}, \qquad \textit{for all } m \in \mathcal{M}.$$

Then the unique equilibrium of the potentialized game coincides with p^*.

Near-Optimal Dynamics

- We will study the dynamic properties of the power game $\mathcal{G}$ when the prices are set equal to λ^*.
- A natural class of dynamics is the best-response dynamics, in which each user updates his strategy to maximize its utility, given the strategies of other users.
- Let $\beta_m : \mathcal{P}_{-m} \to \mathcal{P}_m$ denote the best-response mapping of user m, i.e.,

$$\beta_m(\mathbf{p}_{-m}) = \arg\max_{p_m \in \mathcal{P}_m} u_m(p_m, \mathbf{p}_{-m}).$$

- Discrete time BR dynamics:

$$p_m \leftarrow p_m + \alpha\left(\beta_m(\mathbf{p}_{-m}) - p_m\right) \quad \text{for all } m \in \mathcal{M},$$

- Continuous time BR dynamics:

$$\dot{p}_m = \beta_m(\mathbf{p}_{-m}) - p_m \quad \text{for all } m \in \mathcal{M}.$$

- The continuous-time BR dynamics is similar to continuous time fictitious play dynamics and gradient-play dynamics [Flam, 2002], [Shamma and Arslan, 2005], [Fudenberg and Levine, 1998].

Convergence Analysis – 1

- If users use BR dynamics in the potentialized game $\tilde{\mathcal{G}}$, their strategies converge to the desired operating point p^*.
  - This can be shown through a Lyapunov analysis using the potential function of $\tilde{\mathcal{G}}$, [Hofbauer and Sandholm, 2000]
  - Our interest is in studying the convergence properties of BR dynamics when used in the power game $\mathcal{G}$.
- Idea: Use perturbation analysis from system theory
  - The difference between the utilities of the original and the potentialized game can be viewed as a perturbation.
  - Lyapunov function of the potentialized game can be used to characterize the set to which the BR dynamics for the original power game converges.

Convergence Analysis – 2

- Our first result shows BR dynamics applied to game $\mathcal{G}$ converges to the set of ϵ-equilibria of the potentialized game $\tilde{\mathcal{G}}$, denoted by $\tilde{\mathcal{I}}_\epsilon$.
- We define the minimum SINR:

$$\underline{SINR}_m = \frac{\underline{P}_m h_{mm}}{N_0 + \sum_{k \neq m} h_{km} \overline{P}_k}$$

- We say that the dynamics *converges uniformly* to a set S if there exists some $T \in (0, \infty)$ such that $\mathbf{p}^t \in S$ for every $t \geq T$ and any initial operating point $\mathbf{p}^0 \in \mathcal{P}$.

Lemma

The BR dynamics applied to the original power game $\mathcal{G}$ converges uniformly to the set $\tilde{\mathcal{I}}_\epsilon$, where ϵ satisfies

$$\epsilon \leq \frac{1}{\gamma} \sum_{m \in \mathcal{M}} \frac{1}{\underline{SINR}_m}.$$

- The error bound provides the explicit dependence on γ and $\underline{SINR}_m$.

Convergence Analysis – 3

- We next establish how "far" the power allocations in $\tilde{\mathcal{I}}_\epsilon$ can be from the desired operating point $\mathbf{p}^*$.

Theorem

For all ϵ, $\mathbf{p} \in \tilde{\mathcal{I}}_\epsilon$ satisfies

$$|\tilde{p}_m - p_m^*| \leq \overline{P}_m \sqrt{2\epsilon} \quad \text{for every } \tilde{p} \in \tilde{\mathcal{I}}_\epsilon \text{ and every } m \in \mathcal{M}$$

- Idea: Using the strict concavity and the additively separable structure of the potential function, we characterize $\tilde{\mathcal{I}}_\epsilon$.

Convergence and the System Utility

- Under some smoothness assumptions, the error bound enables us to characterize the performance loss in terms of system utility.

Theorem

Let $\epsilon > 0$ be given. (i) Assume that U_0 is a Lipschitz continuous function, with a Lipschitz constant given by L. Then

$$|U_0(\mathbf{p}^*) - U_0(\tilde{\mathbf{p}})| \leq \sqrt{2\epsilon} L \sqrt{\sum_{m \in \mathcal{M}} \overline{P}_m^2}, \quad \text{for every } \tilde{\mathbf{p}} \in \tilde{\mathcal{I}}_\epsilon.$$

(ii) Assume that U_0 is a continuously differentiable function so that $|\frac{\partial U_0}{\partial p_m}| \leq L_m$, $m \in \mathcal{M}$. Then

$$|U_0(\mathbf{p}^*) - U_0(\tilde{\mathbf{p}})| \leq \sqrt{2\epsilon} \sum_{m \in \mathcal{M}} \overline{P}_m L_m, \quad \text{for every } \tilde{\mathbf{p}} \in \tilde{\mathcal{I}}_\epsilon.$$

Numerical Example – 1

- Consider a system with 3 users and let the desired operating point be given by $\mathbf{p}^* = [5, 5, 5]$.
- We choose the prices as $\lambda_m^* = \frac{1}{p_M^*}$ and pick the channel gain coefficients uniformly at random.
- We consider three different values of $\gamma \in \{5, 10, 50\}$.

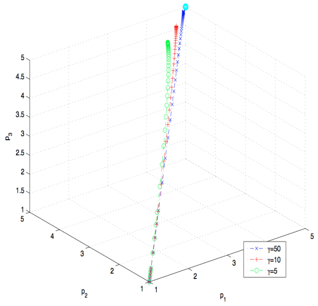

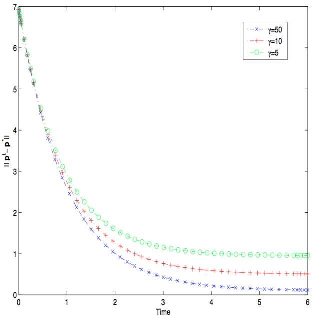

Sum-rate Objective

- We next consider the natural system objective of maximizing the sum-rate in the network.

$$U_0(\mathbf{p}) = \sum_m r_m(\mathbf{p}).$$

- The performance loss in our pricing scheme can be quantified as follows.

Theorem

Let $\mathbf{p}^*$ be the operating point that maximizes sum-rate objective, and let $\tilde{\mathcal{I}}_\epsilon$ be the set of ϵ-equilibria of the modified game to which the BR dynamics converges. Then

$$|U_0(\mathbf{p}^*) - U_0(\tilde{\mathbf{p}})| \leq \sqrt{2\epsilon}(M-1) \sum_{m \in \mathcal{M}} \frac{\overline{P}_m}{\underline{P}_m}, \qquad \text{for every } \tilde{\mathbf{p}} \in \tilde{\mathcal{I}}_\epsilon.$$

Numerical Example – 2

- Consider $M = 10$ users and assume that the power bounds are given by $\underline{P}_m = 10^{-2}$, $\overline{P}_m = 10$ for all $m \in \mathcal{M}$.

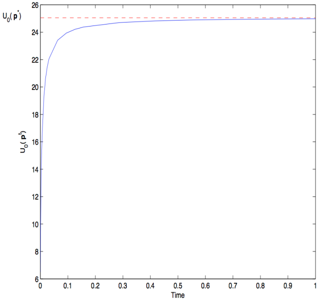

(c) The change in sum-rate as a function of time for $\gamma = 10$.

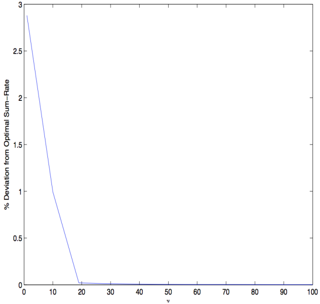

(d) The effect of γ on performance loss.

Summary

- Analysis of the global structure of preferences
- Decomposition: nonstrategic, potential and harmonic components
- Projection to "closest" potential game
- Preserves ϵ-approximate equilibria and dynamics
- Enables extension of many tools to non-potential games

Want to know more?

- Candogan, Menache, Ozdaglar, P., Flow representations of games: harmonic and potential games. Preprint.
- Candogan, Menache, Ozdaglar, P., Near-optimal power control in wireless networks: a potential game approach. INFOCOM 2010.