Network Games: Learning and Dynamics

Asu Ozdaglar

Conference on Decision and Control (CDC)

December 2008

Department of Electrical Engineering & Computer Science

Massachusetts Institute of Technology, USA

Introduction

- **Central Question in Today's and Future Networks:** Systematic analysis and design of network architectures and development of network control schemes
- **Traditional Network Optimization:** Single administrative domain with a single control objective and obedient users.
- **New Challenges:**
  - Large-scale and interconnection of heterogeneous autonomous entities
    - Control in the presence of selfish incentives and private information of users
  - Continuous upgrades and investments in new technologies
    - Economic incentives of service and content providers more paramount
  - New situation-aware wireless technologies to deal with inherent dynamics
    - Autonomous decisions based on current network conditions
  - Analysis of social and economic networks
    - Learning, information aggregation, control, endogenous network formation
- These challenges make **game theory and economic market mechanisms** natural tools for the analysis of large-scale networked systems

Issues in Network Games

- Game theory has traditionally been used in economics and social sciences with focus on fully rational interactions
  - Theory developed for small scale sophisticated interactions
  - **Strong assumptions:** common knowledge, common prior, forward-looking behavior
- In (engineering or social) networked systems, not necessarily a good framework for two reasons:
  - Large-scale systems consisting of individuals with partial information
  - Most focus on dynamic interactions and in particular **learning dynamics**

Learning Dynamics in Games

- **Bayesian Learning:**
  - Update beliefs (about an underlying state or opponent strategies) based on new information optimally (i.e., in a Bayesian manner)
- **Adaptive Learning:**
  - Myopic, simple and rule-of-thumb
  - **Example:** Fictitious play
    - Play optimally against the empirical distribution of past play of opponent
- **Evolutionary Dynamics:**
  - Selection of strategies according to performance against aggregates and random mutations

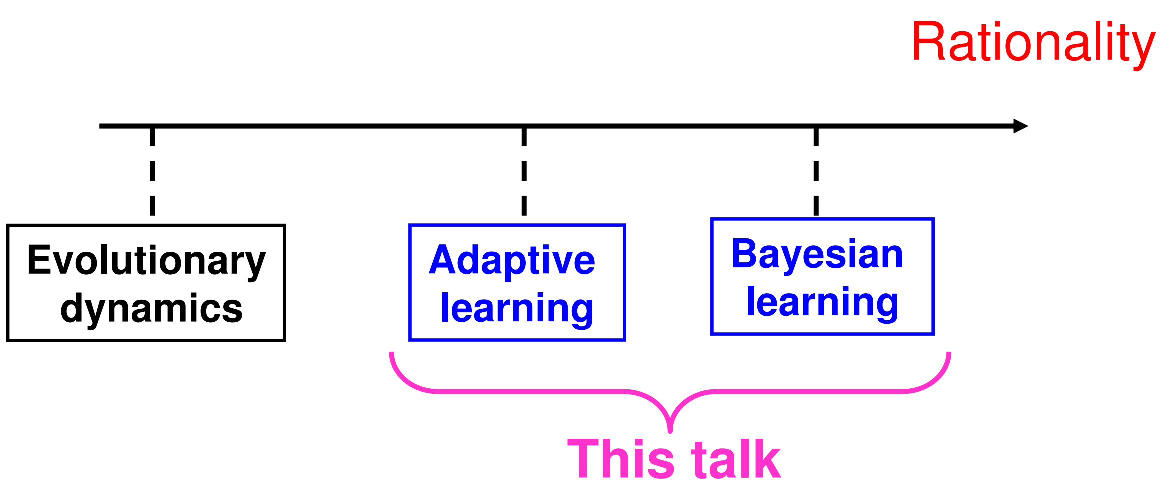

This Tutorial

- Strategic form games and Nash equilibrium
- Adaptive learning in games
  - Fictitious play and shortcomings
- Special classes of games:
  - Supermodular games and dynamics
  - Potential and congestion games and dynamics
- Bayesian learning in games
  - Information aggregation in social networks

Strategic Form Games

- A strategic (form) game is a model for a game in which all of the participants act simultaneously and without knowledge of other players' actions.

Definition (Strategic Game): A *strategic game* is a triplet $\langle \mathcal{I}, (S_i)_{i\in\mathcal{I}}, (u_i)_{i\in\mathcal{I}} \rangle$:

- $\mathcal{I}$ is a finite set of players, $\mathcal{I} = \{1, \ldots, I\}$.
- S_i is the set of available actions for player i
  - $s_i \in S_i$ is an action for player i
  - $s_{-i} = [s_j]_{j\neq i}$ is a vector of actions for all players *except* i.
  - $(s_i, s_{-i}) \in S$ is an *action profile*, or *outcome*.
  - $S = \prod_i S_i$ is the set of all action profiles
  - $S_{-i} = \prod_{j\neq i} S_j$ is the set of all action profiles for all players *except* i
- $u_i : S \to \mathbb{R}$ is the payoff (utility) function of player i

- We will use the terms action and pure strategy interchangeably.

Example

- **Example:** Cournot competition.
  - Two firms producing the same good.
  - The action of a player i is a quantity, $s_i \in [0, \infty]$ (amount of good he produces).
  - The utility for each player is its total revenue minus its total cost,

$$u_i(s_1, s_2) = s_i p(s_1 + s_2) - c s_i$$

    where $p(q)$ is the price of the good (as a function of the total amount), and c is unit cost (same for both firms).

- Assume for simplicity that $c = 1$ and $p(q) = \max\{0, 2 - q\}$
- Consider the best-response correspondences for each of the firms, i.e., for each i, the mapping $B_i(s_{-i}) : S_{-i} \to S_i$ such that

$$B_i(s_{-i}) \in \mathrm{argmax}_{s_i \in S_i} u_i(s_i, s_{-i}).$$

Example–Continued

- By using the first order optimality conditions, we have

$$B_i(s_{-i}) = \text{argmax}_{s_i \geq 0}(s_i(2 - s_i - s_{-i}) - s_i)$$
$$= \begin{cases} \frac{1-s_{-i}}{2} & \text{if } s_{-i} \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

- The figure illustrates the best response functions as a function of s_1 and s_2.

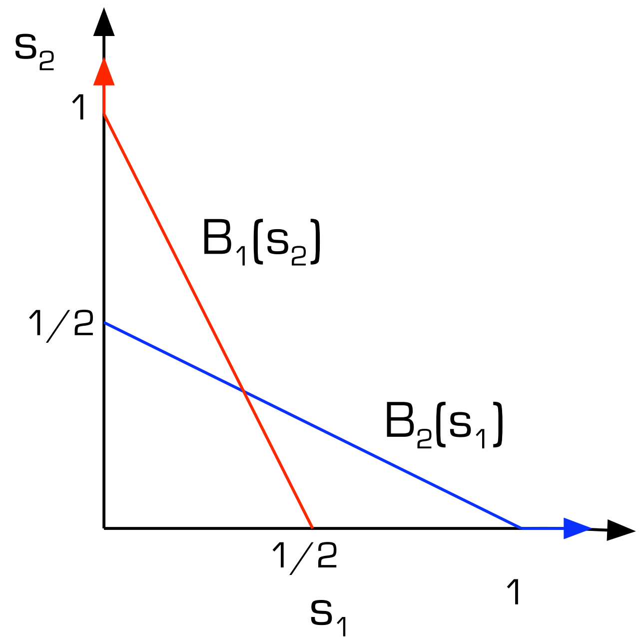

- Assuming that players are rational and fully knowledgable about the structure of the game and each other's rationality, what should the outcome of the game be?

Pure and Mixed Strategy Nash Equilibrium

Definition (Nash equilibrium): A (pure strategy) Nash Equilibrium of a strategic game $\langle \mathcal{I}, (S_i)_{i\in\mathcal{I}}, (u_i)_{i\in\mathcal{I}} \rangle$ is a strategy profile $s^* \in S$ such that for all $i \in \mathcal{I}$

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \qquad \text{for all } s_i \in S_i.$$

- No player can profitably deviate given the strategies of the other players
- An action profile s^* is a Nash equilibrium if and only if

$$s_i^* \in B_i(s_{-i}^*) \qquad \text{for all } i \in \mathcal{I},$$

- Let Σ_i denote the set of probability measures over the pure strategy set S_i.
- We use $\sigma_i \in \Sigma_i$ to denote the **mixed strategy of player** i, and $\sigma \in \Sigma = \prod_{i\in\mathcal{I}} \Sigma_i$ to denote a **mixed strategy profile** (similarly define $\sigma_{-i} \in \Sigma_{-i} = \prod_{j\neq i} \Sigma_j$)
- Following Von Neumann-Morgenstern expected utility theory, we extend the payoff functions u_i from S to Σ by

$$u_i(\sigma) = \int_S u_i(s) d\sigma(s).$$

Definition (Mixed Nash Equilibrium): A mixed strategy profile σ^* is a (mixed strategy) Nash Equilibrium if for each player i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \qquad \text{for all } \sigma_i \in \Sigma_i.$$

Existence of Nash Equilibria

Theorem: [Nash 50] Every finite game has a mixed strategy Nash equilibrium.

Proof Outline:

- σ^* mixed Nash equilibrium if and only if $\sigma_i^* \in B_i(\sigma_{-i}^*)$ for all $i \in \mathcal{I}$, where

$$B_i(\sigma_{-i}^*) \in \arg\max_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \sigma_{-i}^*).$$

- This can be written compactly as $\sigma^* \in B(\sigma^*)$, where $B(\sigma) = [B_i(\sigma_{-i})]_{i \in \mathcal{I}}$, i.e., σ^* is a **fixed point of the best-response correspondence**.
- Use Kakutani's fixed point theorem to establish the existence of a fixed point.

Linearity of expectation in probabilities play a key role; extends to (quasi)-concave payoffs in infinite games

Theorem: [Debreu, Glicksberg, Fan 52] Assume that the S_i are nonempty compact convex subsets of an Euclidean space. Assume that the payoff functions $u_i(s_i, s_{-i})$ are quasi-concave in s_i and continuous in s, then there exists a pure strategy Nash equilibrium.

- Existence of mixed strategy equilibria for continuous games [Glicksberg 52] and some discontinuous games [Dasgupta and Maskin 86]

Adaptive Learning in Finite Games

- Most economic theory relies on equilibrium analysis based on Nash equilibrium or its refinements.
- **Traditional explanation for when and why equilibrium arises:**
  - Results from analysis and introspection by sophisticated players when the structure of the game and the rationality of the players are all common knowledge.
- **Alternative justification more relevant for networked-systems:**
  - Arises as the limit point of a repeated play in which less than fully rational players myopically update their behavior
  - Agents behave as if facing a stationary, but unknown, distribution of opponents' strategies

Fictitious Play

- A natural and widely used model of learning is fictitious play [Brown 51]
  - Players form beliefs about opponent play and myopically optimize their action with respect to these beliefs
- Agent i forms the empirical frequency distribution of his opponent j's past play according to

$$\mu_j^t(\tilde{s}_j) = \frac{1}{t} \sum_{\tau=0}^{t-1} I(s_j^t = \tilde{s}_j),$$

  let $\mu_{-i}^t = \prod_{j \neq i} \mu_j^t$ for all t.

- He then chooses his action at time t to maximize his payoff, i.e.,

$$s_i^t \in \arg\max_{s_i \in S_i} u_i(s_i, \mu_{-i}^t).$$

  - This choice is myopic, i.e., players are trying to maximize current payoff without considering their future payoffs.
  - Players only need to know their own utility function.

Basic Properties of Fictitious Play

- Let $\{s^t\}$ be a sequence of strategy profiles generated by fictitious play.
- We say that $\{s^t\}$ converges to $\sigma \in \Sigma$ in the time-average sense if the empirical frequencies converge to σ, i.e., $\mu_i^t \rightarrow \sigma_i$ for all i.

Proposition: Suppose a fictitious play sequence $\{s^t\}$ converges to σ in the time-average sense. Then σ is a Nash equilibrium of the stage game.

- Is convergence in the time-average sense a natural notion of convergence?

Shortcomings of Fictitious Play

Mis-coordination example [Fudenberg, Kreps 88]: Consider the FP of the game:

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Note that this game had a unique mixed Nash equilibrium $\big((1/2, 1/2), (1/2, 1/2)\big)$.
Consider the following sequence of play:

	η_1^t	η_2^t	Play
0	(0,1/2)	(1/2,0)	(A,B)
1	(1,1/2)	(1/2,1)	(B,A)
2	(1,3/2)	(3/2,1)	(A,B)
3	$\cdots$	$\cdots$	(B,A)

- Play continues as (A,B), (B,A), ... - a deterministic cycle.
- The time average converges to $\big((1/2, 1/2), (1/2, 1/2)\big)$, which is a mixed strategy equilibrium of the game.
- But players never successfully coordinate!

Alternative Focus

- Various convergence problems present for adaptive learning rules
  - Uncoupled dynamics do not lead to Nash equilibrium! [Hart, Mas-Colell 03]
- Rather than seeking learning dynamics that converge to reasonable behavior in all games, focus on relevant classes games that arise in engineering and economics
- **In particular, this talk:**
  - Supermodular Games
  - Potential Games
- **Advantages:**
  - Tractable and elegant characterization of equilibria, sensitivity analysis
  - Most reasonable adaptive learning rules converge to Nash equilibria

Supermodular Games

- Supermodular games are those characterized by **strategic complementarities**
- Informally, this means that the **marginal utility of increasing a player's strategy raises with increases in the other players' strategies**.
- **Why interesting?**
  - They arise in many models.
  - Existence of a pure strategy equilibrium without requiring the quasi-concavity of the payoff functions.
  - Many solution concepts yield the same predictions.
  - The equilibrium set has a smallest and a largest element.
  - They have nice sensitivity (or comparative statics) properties and behave well under a variety of distributed dynamic rules.
- The machinery needed to study supermodular games is lattice theory and monotonicity results in lattice programming
  - Methods used are **non-topological and they exploit order properties**

Increasing Differences

- We first study the monotonicity properties of optimal solutions of parametric optimization problems:

$$x(t) \in \arg\max_{x \in X} f(x, t),$$

where $f : X \times T \to \mathbb{R}$, $X \subset \mathbb{R}$, and T is some partially ordered set.

Definition: Let $X \subseteq \mathbb{R}$ and T be some partially ordered set. A function $f : X \times T \to \mathbb{R}$ has **increasing differences** in (x, t) if for all $x' \geq x$ and $t' \geq t$, we have

$$f(x', t') - f(x, t') \geq f(x', t) - f(x, t).$$

- incremental gain to choosing a higher x (i.e., x' rather than x) is greater when t is higher, i.e., $f(x', t) - f(x, t)$ is nondecreasing in t.

Lemma: Let $X \subseteq \mathbb{R}$ and $T \subset \mathbb{R}^k$ for some k, a partially ordered set with the usual vector order. Let $f : X \times T \to \mathbb{R}$ be a twice continuously differentiable function. Then, the following statements are equivalent:

(a) The function f has increasing differences in (x, t).

(b) For all $x \in X$, $t \in T$, and all $i = 1, \ldots, k$, we have

$$\frac{\partial^2 f(x, t)}{\partial x \partial t_i} \geq 0.$$

Examples–I

Example: Network effects (positive externalities).

- A set $\mathcal{I}$ of users can use one of two technologies X and Y (e.g., Blu-ray and HD DVD)
- $B_i(J, k)$ denotes payoff to i when a subset J of users use technology k and $i \in J$
- There exists a network effect or positive externality if

$$B_i(J, k) \leq B_i(J', k), \qquad \text{when } J \subset J',$$

  i.e., player i better off if more users use the same technology as him.
- Leads naturally to a strategic form game with actions $S_i = \{X, Y\}$
- Define the order $Y \succeq X$, which induces a lattice structure
- Given $s \in S$, let $X(s) = \{i \in \mathcal{I} \mid s_i = X\}$, $Y(s) = \{i \in \mathcal{I} \mid s_i = Y\}$.
- Define the payoffs as

$$u_i(s_i, s_{-i}) = \begin{cases} B_i(X(s), X) & \text{if } s_i = X, \\ B_i(Y(s), Y) & \text{if } s_i = Y \end{cases}$$

- Show that the payoff functions of this game feature increasing differences.

Examples –II

Example: Cournot duopoly model.

- Two firms choose the quantity they produce $q_i \in [0, \infty)$.
- Let $P(Q)$ with $Q = q_i + q_j$ denote the inverse demand (price) function. Payoff function of each firm is $u_i(q_i, q_j) = q_i P(q_i + q_j) - cq_i$.
- Assume $P'(Q) + q_i P''(Q) \leq 0$ (firm i's marginal revenue decreasing in q_j).
- Show that the payoff functions of the transformed game defined by $s_1 = q_1$, $s_2 = -q_2$ has increasing differences in (s_1, s_2).

Monotonicity of Optimal Solutions

Theorem: [Topkis 79] Let $X \subset \mathbb{R}$ be a compact set and T be some partially ordered set. Assume that the function $f : X \times T \to \mathbb{R}$ is upper semicontinuous in x for all $t \in T$ and has increasing differences in (x, t). Define $x(t) = \arg\max_{x \in X} f(x, t)$. Then, we have:

1. For all $t \in T$, $x(t)$ is nonempty and has a greatest and least element, denoted by $\bar{x}(t)$ and $\underline{x}(t)$ respectively.

2. For all $t' \geq t$, we have $\bar{x}(t') \geq \bar{x}(t)$ and $\underline{x}(t') \geq \underline{x}(t)$.

- If f has increasing differences, the set of optimal solutions $x(t)$ is non-decreasing in the sense that the largest and the smallest selections are non-decreasing.

Supermodular Games

Definition: The strategic game $\langle \mathcal{I}, (S_i), (u_i) \rangle$ is a supermodular game if for all i:

1. S_i is a compact subset of $\mathbb{R}$ (or more generally S_i is a complete lattice in $\mathbb{R}^{m_i}$),
2. u_i is upper semicontinuous in s_i, continuous in s_{-i},
3. u_i has increasing differences in (s_i, s_{-i}) [or more generally u_i is supermodular in (s_i, s_{-i}), which is an extension of the property of increasing differences to games with multi-dimensional strategy spaces].

- Apply Topkis' Theorem to best response correspondences

Corollary: Assume $\langle \mathcal{I}, (S_i), (u_i) \rangle$ is a supermodular game. Let

$$B_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i}).$$

Then:

1. $B_i(s_{-i})$ has a greatest and least element, denoted by $\bar{B}_i(s_{-i})$ and $\underline{B}_i(s_{-i})$.
2. If $s'_{-i} \geq s_{-i}$, then $\bar{B}_i(s'_{-i}) \geq \bar{B}_i(s_{-i})$ and $\underline{B}_i(s'_{-i}) \geq \underline{B}_i(s_{-i})$.

Existence of a Pure Nash Equilibrium

- Follows from Tarski's fixed point theorem

Theorem: [Tarski 55] Let S be a compact sublattice of $\mathbb{R}^k$ and $f : S \to S$ be an increasing function (i.e., $f(x) \leq f(y)$ if $x \leq y$). Then, the set of fixed points of f, denoted by E, is nonempty.

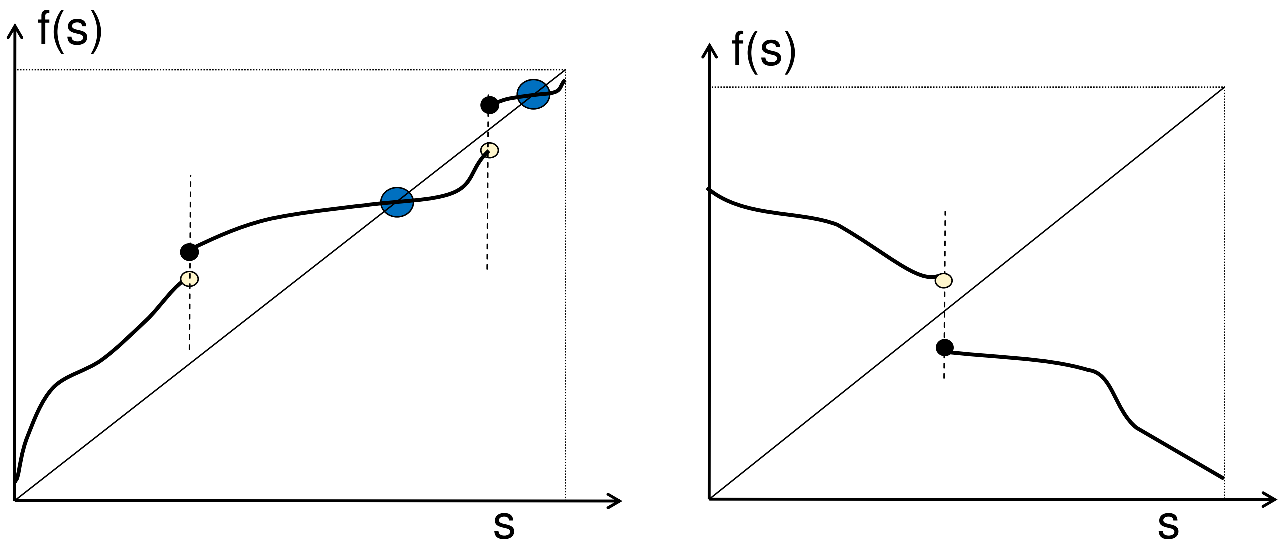

- Apply Tarski's fixed point theorem to best response correspondences
- Nash equilibrium set has a largest and a smallest element, and easy sensitivity results (e.g., quantity supplied increases with demand in Cournot game)

Dynamics in Supermodular Games

Theorem: [Milgrom, Roberts 90] Let $G = \langle \mathcal{I}, (S_i), (u_i) \rangle$ be a supermodular game. Let $\{s^t\}$ be a sequence of strategy profiles generated by **reasonable adaptive learning rules**. Then,

$$\liminf_{t\to\infty} s^t \geq \underline{s} \quad \text{and} \quad \limsup_{t\to\infty} s^t \leq \bar{s},$$

where $\underline{s}$ and $\bar{s}$ are smallest and largest Nash equilibria of G.

Reasonable adaptive learning rules: Best-response, fictitious play ...

Remarks:

- Implies convergence for games with unique Nash equilibrium.
- Fictitious play converges for general supermodular games [Krishna 92], [Berger 03, 07], [Hahn 08]

Example: Apply best-response dynamics to Cournot game

Wireless Power Control Game

- Power control in cellular CDMA wireless networks [Alpcan, Basar, Srikant, Altman 02], [Gunturi, Paganini 03]

- It has been recognized that in the presence of interference, the strategic interactions between the users is that of **strategic complementarities** [Saraydar, Mandayam, Goodman 02], [Altman and Altman 03]

Model:

- Let $L = \{1, 2, ..., n\}$ denote the set of users (nodes) and $\mathcal{P} = \prod_{i \in L} [P_i^{\text{min}}, P_i^{\text{max}}]$ denote the set of power vectors $p = [p_1, \ldots, p_n]$.

- Each user is endowed with a utility function $f_i(\gamma_i)$ as a function of its SINR γ_i.
  - $f_i(\gamma_i)$ depends on details of transmission: modulation, coding, packet size
  - In most practical cases, $f(\gamma)$ is strictly increasing and has a sigmoidal shape.

- The payoff function of each user represents a tradeoff between the payoff obtained by the received SINR and the power expenditure, and takes the form

$$u_i(p_i, p_{-i}) = f_i(\gamma_i) - cp_i.$$

Increasing Differences

- Assume that each utility function satisfies the following assumption regarding its **coefficient of relative risk aversion**:

$$\frac{-\gamma_i f_i''(\gamma_i)}{f_i'(\gamma_i)} \geq 1, \qquad \text{for all } \gamma_i \geq 0.$$

  - Satisfied by α-fair functions $f(\gamma) = \frac{\gamma^{1-\alpha}}{1-\alpha}$, $\alpha > 1$ [Mo, Walrand 00], and the efficiency functions introduced earlier

- Show that for all i, the function $u_i(p_i, p_{-i})$ has increasing differences in (p_i, p_{-i}).

Implications:

- Power control game has a pure Nash equilibrium.
- The Nash equilibrium set has a largest and a smallest element, and there are **distributed algorithms** that will converge to any of these equilibria.
- These algorithms involve each user updating their power level locally (based on total received power at the base station).

Potential Games

Definition [Monderer and Shapley 96]:

(i) A function $\Phi : S \to \mathbb{R}$ is called an ordinal potential function for the game G if for all i and all $s_{-i} \in S_{-i}$,

$$u_i(x, s_{-i}) - u_i(z, s_{-i}) > 0 \quad \text{iff} \quad \Phi(x, s_{-i}) - \Phi(z, s_{-i}) > 0, \quad \text{for all } x, z \in S_i.$$

(ii) A function $\Phi : S \to \mathbb{R}$ is called a potential function for the game G if for all i and all $s_{-i} \in S_{-i}$,

$$u_i(x, s_{-i}) - u_i(z, s_{-i}) = \Phi(x, s_{-i}) - \Phi(z, s_{-i}), \quad \text{for all } x, z \in S_i.$$

G is called an ordinal (exact) potential game if it admits an ordinal (exact) potential.

Properties of Potential Games

- A global maximum of an ordinal potential function is a pure Nash equilibrium (there may be other pure NE, which are local maxima)
  - Every finite ordinal potential game has a pure Nash equilibrium.
- Many adaptive learning dynamics "converge" to a pure Nash equilibrium [Monderer and Shapley 96], [Young 98, 05], [Hart, Mas-Colell 00,03], [Marden, Arslan, Shamma 06, 07]
  - **Examples:** Fictitious play, better reply with inertia, spatial adaptive play, regret matching (for 2 player potential games)

Congestion Games

- Congestion games arise when users need to share resources in order to complete certain tasks
  - For example, drivers share roads, each seeking a minimal cost path.
  - The cost of each road segment adversely affected by the number of other drivers using it.
- **Congestion Model:** $C = \langle N, M, (S_i)_{i \in N}, (c^j)_{j \in M} \rangle$ where
  - $N = \{1, 2, \cdots, n\}$ is the set of players,
  - $M = \{1, 2, \cdots, m\}$ is the set of resources,
  - S_i consists of sets of resources (e.g., paths) that player i can take.
  - $c^j(k)$ is the cost to each user who uses resource j if k users are using it.
- Define congestion game $\langle N, (S_i), (u_i) \rangle$ with utilities $u_i(s_i, s_{-i}) = \sum_{j \in s_i} c^j(k_j)$, where k_j is the number of users of resource j under strategies s.

Theorem: [Rosenthal 73] Every congestion game is a potential game.

Proof idea: Verify that the following is a potential function for the congestion game:

$$\Phi(s) = \sum_{j \in \cup s_i} \left(\sum_{k=1}^{k_j} c^j(k) \right)$$

Network Design

- Sharing the cost of a designed network among participants [Anshelevich et al. 05]

Model:

- Directed graph $N = (V, E)$ with edge cost $c_e \geq 0$, k players
- Each player i has a set of nodes T_i he wants to connect
- A strategy of player i set of edges $S_i \subset E$ such that S_i connects to all nodes in T_i

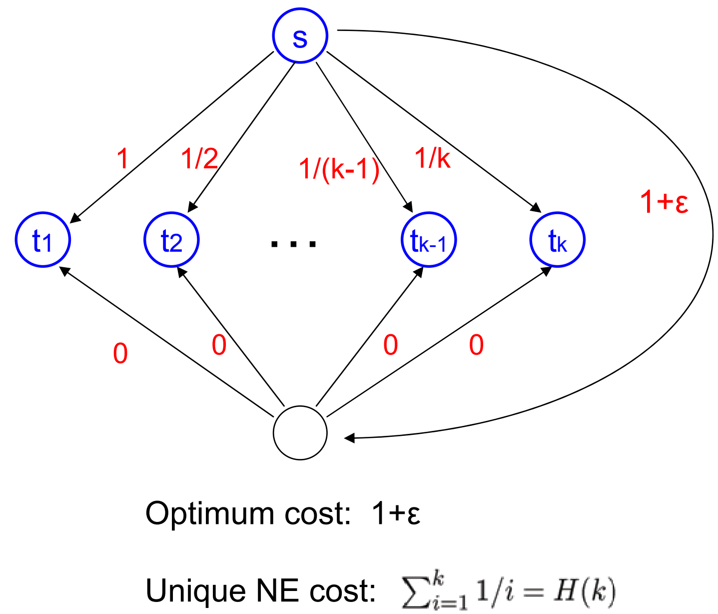

- **Cost sharing mechanism:** All players using an edge split the cost equally
- Given a vector of player's strategies $S = (S_1, \ldots, S_k)$, the cost to agent i is $C_i(S) = \sum_{e \in S_i}(c_e/x_e)$, where x_e is the number of agents whose strategy contains edge e

This game is a **congestion game**, implying existence of a pure Nash equilibrium and convergence of learning dynamics.

Other Examples

Game Theory for Nonconvex Distributed Optimization:

- Distributed Power Control for Wireless Adhoc Networks [Huang,Berry,Honig 05]
  - Two models: Single channel spread spectrum, Multi-channel orthogonal frequency division multiplexing
  - Asynchronous distributed algorithm for optimizing total network performance
  - Convergence analysis in the presence of nonconvexities using **supermodular game theory**
- Distributed Cooperative Control–"Constrained Consensus" [Marden, Arslan, Shamma 07]
  - Distributed algorithms to reach consensus in the "values of multiple agents" (e.g. averaging and rendezvous problems)
  - Nonconvex constraints in agent values
  - Design a game (i.e., utility functions of players) such that
    - The resulting game is a **potential game** and the Nash equilibrium "coincides" with the social optimum
    - Use learning dynamics for potential games to design distributed algorithms with favorable convergence properties

Bayesian Learning in Games

- So far focus on adaptive learning
- Individuals do not update their model even tough they repeatedly observe the strategies of their opponents changing dynamically
- **Alternative paradigm:** Individuals engage in Bayesian updating with (some) understanding of the strategy profiles of others
  - Similar to Bayesian learning in decision-theoretic problems, though richer because of strategic interactions

Model of Bayesian Learning

- Illustrate main issues with a simple model in which learning is about payoff relevant state of the world
- **Relevance to networks:** Model society, information flows as a **social network**
- Dynamic game with sequential decisions based on private signals and observation of past actions
- Payoffs conditional on the (unknown) state of the world
- **Measure of information aggregation:** whether there will be convergence to correct beliefs and decisions in large networks—**asymptotic learning**
- **Question:** Under what conditions—structure of signals, network/communication structure, heterogeneity of preferences—do individuals learn the state as the social network grows bigger?

Difficulties of Bayesian Learning in Games

- Model for Bayesian learning on a line [Bikchandani, Hirschleifer, Welch (92), Banerjee (92)]
- Two possible states of the world $\theta \in \{0,1\}$, both equally likely
- A sequence of agents $(n = 1, 2, ...)$ making decisions $x_n \in \{0,1\}$
- Agent n obtains utility 1 if $x_n = \theta$ and utility 0 otherwise
- Each agent has iid private binary signals s_n, where $s_n = \theta$ with probability $> 1/2$
- Agent n knows his signal s_n and the decisions of previous agents $x_1, x_2, ..., x_{n-1}$
- Agent n chooses action 1 if

$$\mathbb{P}\left(\theta = 1 | s_n, x_1, x_2, ..., x_{n-1}\right) > \mathbb{P}\left(\theta = 0 | s_n, x_1, x_2, ..., x_{n-1}\right)$$

- If $s_1 = s_2 \neq \theta$, then all agents **herd** and $x_n \neq \theta$ for all agents,

$$\lim_{n\to\infty} \mathbb{P}(x_n = \theta) < 1$$

Bayesian Learning in Networks

- Model of learning on networks [Acemoglu, Dahleh, Lobel, Ozdaglar 08]
- Two possible states of the world $\theta \in \{0, 1\}$, both equally likely,
- A sequence of agents $(n = 1, 2, ...)$ making decisions $x_n \in \{0, 1\}$.
- Agent n obtains utility 1 if $x_n = \theta$ and utility 0 otherwise
- Each agent has an iid private signal s_n in S. The signal is generated according to distribution $\mathbb{F}_\theta$, $\mathbb{F}_0$ and $\mathbb{F}_1$ absolutely continuous with respect to each other
- $(\mathbb{F}_0, \mathbb{F}_1)$ is the signal structure
- Agent n has a neighborhood $B(n) \subseteq \{1, 2, ..., n-1\}$ and observes the decisions x_k for all $k \in B(n)$. The set $B(n)$ is private information.
- The neighborhood $B(n)$ is generated according to an arbitrary distribution $\mathbb{Q}_n$
- $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is the network topology and is common knowledge
- A social network consists of the signal structure and network topology
- **Asymptotic Learning:** Under what conditions does $\lim_{n\to\infty} \mathbb{P}(x_n = \theta) = 1$?

Perfect Bayesian Equilibria

- Agent n's information set is $\mathcal{I}_n = \{s_n, B(n), x_k \text{ for all } k \in B(n)\}$
- A strategy for individual n is $\sigma_n : \mathcal{I}_n \to \{0, 1\}$
- A strategy profile is a sequence of strategies $\sigma = \{\sigma_n\}_{n \in \mathbb{N}}$.
  - A strategy profile σ induces a probability measure $\mathbb{P}_\sigma$ over $\{x_n\}_{n \in \mathbb{N}}$.

Definition: A strategy profile σ^* is a pure-strategy **Perfect Bayesian Equilibrium** if for each $n \in \mathbb{N}$

$$\sigma_n^*(\mathcal{I}_n) \in \text{argmax}_{y \in \{0,1\}} \mathbb{P}_{(y, \sigma^*_{-n})}(y = \theta \ |\mathcal{I}_n)$$

- A pure strategy PBE exists. Denote the set of PBEs by Σ^*.

Definition: Given a signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ and a network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$, we say that **asymptotic learning occurs in equilibrium** σ if x_n converges to θ in probability (according to measure $\mathbb{P}_\sigma$), that is,

$$\lim_{n \to \infty} \mathbb{P}_\sigma(x_n = \theta) = 1$$

Equilibrium Decision Rule

Lemma: The decision of agent n, $x_n = \sigma(\mathcal{I}_n)$, satisfies

$$x_n = \begin{cases} 1, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid s_n) + \mathbb{P}_\sigma\big(\theta = 1 \mid B(n), x_k \text{ for all } k \in B(n)\big) > 1, \\ 0, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid s_n) + \mathbb{P}_\sigma\big(\theta = 1 \mid B(n), x_k \text{ for all } k \in B(n)\big) < 1, \end{cases}$$

and $x_n \in \{0, 1\}$ otherwise.

- **Implication:** The belief about the state decomposes into two parts:
  - the **Private Belief**: $\mathbb{P}_\sigma(\theta = 1 \mid s_n)$;
  - the **Social Belief**: $\mathbb{P}_\sigma(\theta = 1 \mid B(n), x_k \text{ for all } k \in \omega_n)$.

Private Beliefs

Lemma: The private belief of agent n is

$$p_n(s_n) = \mathbb{P}_\sigma(\theta = 1 | s_n) = \left(1 + \frac{d\mathbb{F}_0(s_n)}{d\mathbb{F}_1(s_n)}\right)^{-1}.$$

Definition: The signal structure has **bounded private beliefs** if there exists some $0 < m, M < \infty$ such that the Radon-Nikodym derivate $d\mathbb{F}_0/d\mathbb{F}_1$ satisfies

$$m < \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s) < M,$$

for almost all $s \in S$ under measure $(\mathbb{F}_0 + \mathbb{F}_1)/2$. The signal structure has **unbounded private beliefs** if

$$\inf_{s \in S} \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s) = 0 \quad \text{and} \quad \sup_{s \in S} \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s) = \infty.$$

- Bounded private beliefs $\Leftrightarrow$ bounded likelihood ratio
- If the private beliefs are unbounded, then there exist some agents with **beliefs arbitrarily close to 0** and other agents with **beliefs arbitrarily close to 1**.

Properties of Network Topology

Definition: A network topology $\{\mathbb{Q}_n\}_{n\in\mathbb{N}}$ has **expanding observations** if for all K,

$$\lim_{n\to\infty} \mathbb{Q}_n\left(\max_{b\in B(n)} b < K\right) = 0.$$

Otherwise, it has **nonexpanding observations**

- Expanding observations do not imply connected graph
- Nonexpanding observations equivalently : There exists some K, $\epsilon > 0$ and an infinite subset $\mathcal{N} \in \mathbb{N}$ such that

$$\mathbb{Q}_n\left(\max_{b\in B(n)} b < K\right) \geq \epsilon \quad \text{for all} \quad n \in \mathcal{N}.$$

- A finite group of agents is **excessively influential** if there exists an infinite number of agents who, with probability uniformly bounded away from 0, observe only the actions of a subset of this group.
  - For example, a group is excessively influential if it is the source of *all information* for an infinitely large component of the network
- Nonexpanding observations $\Leftrightarrow$ excessively influential agents

Main Results - I

Theorem 1: Assume that the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has nonexpanding observations. Then, there exists no equilibrium $\sigma \in \Sigma^*$ with asymptotic learning.

Theorem 2: Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has unbounded private beliefs and the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has expanding observations. Then, asymptotic learning occurs in every equilibrium $\sigma \in \Sigma^*$.

- **Implication:** Influential, but **not** excessively influential, individuals (observed by disproportionately more agents in the future) do not prevent learning.
- This contrasts with results in models of myopic learning
- **Intuition:** because the weight given to the information of influential individuals is reduced according to Bayesian updating.

Main Results - II

Theorem 3: If the private beliefs are bounded and the network topology satisfies one of the following conditions,

(a) $B(n) = \{1, ..., n-1\}$ for all n or $|B(n)| \leq 1$ for all n,

(b) there exists some constant M such that $|B(n)| \leq M$ for all n and

$$\lim_{n\to\infty} \max_{b \in B(n)} b = \infty \text{ with probability 1,}$$

then asymptotic learning does not occur.

- **Implication:** No learning with random sampling and bounded beliefs

Theorem 4: There exist network topologies where asymptotic learning occurs for any signal structure $(\mathbb{F}_0, \mathbb{F}_1)$.

Example: For all n,

$$B(n) = \begin{cases} \{1, ..., n-1\}, & \text{with probability } 1 - r(n); \\ \emptyset, & \text{with probability } r(n), \end{cases}$$

for some sequence $\{r(n)\}$ where $\lim_{n\to\infty} r(n) = 0$ and $\sum_{n=1}^{\infty} r(n) = \infty$.

In this case, asymptotic learning occurs for an arbitrary signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ and at any equilibrium.

Concluding Remarks

- Game theory increasingly used for the analysis and control of networked systems
- Many applications:
  - Sensor networks, mobile ad hoc networks
  - Large-scale data networks, Internet
  - Social and economic networks
  - Electricity and energy markets
- Future Challenges
  - Models for understanding when equilibrium behavior yields efficient outcomes
  - Dynamics of agent interactions over large-scale networks
  - Endogenous network formation: dynamics of decisions and graphs
  - Interactions of heterogeneous interlayered networks (e.g., social and communication networks)