6.254 : Game Theory with Engineering Applications
Lecture 6: Continuous and Discontinuous Games

Asu Ozdaglar
MIT

February 23, 2010

Outline

- Continuous Games
- Existence of a Mixed Nash Equilibrium in Continuous Games (Glicksberg's Theorem)
- Existence of a Mixed Nash Equilibrium with Discontinuous Payoffs
- Construction of a Mixed Nash Equilibrium with Infinite Strategy Sets
- Uniqueness of a Pure Nash Equilibrium for Continuous Games

- Reading:
  - Myerson, Chapter 3.
  - Fudenberg and Tirole, Sections 12.2, 12.3.
  - Rosen J.B., "Existence and uniqueness of equilibrium points for concave N-person games," *Econometrica*, vol. 33, no. 3, 1965.

Continuous Games

- We consider games in which players may have infinitely many pure strategies.

Definition

A **continuous game** *is a game* $\langle \mathcal{I}, (S_i), (u_i) \rangle$ *where* $\mathcal{I}$ *is a finite set, the* S_i *are nonempty compact metric spaces, and the* $u_i : S \rightarrow \mathbb{R}$ *are continuous functions.*

- We next state the analogue of Nash's Theorem for continuous games.

Existence of a Mixed Nash Equilibrium

Theorem

(Glicksberg) *Every continuous game has a mixed strategy Nash equilibrium.*

- With continuous strategy spaces, space of mixed strategies infinite dimensional, therefore we need a more powerful fixed point theorem than the version of Kakutani we have used before.
- Here we adopt an alternative approach to prove Glicksberg's Theorem, which can be summarized as follows:
  - We approximate the original game with a sequence of finite games, which correspond to successively finer discretization of the original game.
  - We use Nash's Theorem to produce an equilibrium for each approximation.
  - We use the weak topology and the continuity assumptions to show that these converge to an equilibrium of the original game.

Closeness of Two Games

- Let $u = (u_1, \ldots, u_I)$ and $\tilde{u} = (\tilde{u}_1, \ldots, \tilde{u}_I)$ be two profiles of utility functions defined on S such that for each $i \in \mathcal{I}$, the functions $u_i : S \to \mathbb{R}$ and $\tilde{u}_i : S \to \mathbb{R}$ are bounded (measurable) functions.
- We define the distance between the utility function profiles u and $\tilde{u}$ as

$$\max_{i \in \mathcal{I}} \sup_{s \in S} |u_i(s) - \tilde{u}_i(s)|.$$

- Consider two strategic form games defined by two profiles of utility functions:

$$G = \langle \mathcal{I}, (S_i), (u_i) \rangle, \qquad \tilde{G} = \langle \mathcal{I}, (S_i), (\tilde{u}_i) \rangle.$$

- If σ is a mixed strategy Nash equilibrium of G, then σ need not be a mixed strategy Nash equilibrium of $\tilde{G}$.
- Even if u and $\tilde{u}$ are very close, the equilibria of G and $\tilde{G}$ may be far apart.
  - For example, assume there is only one player, $S_1 = [0, 1]$, $u_1(s_1) = \epsilon s_1$, and $\tilde{u}_1(s_1) = -\epsilon s_1$, where $\epsilon > 0$ is a sufficiently small scalar. The unique equilibrium of G is $s_1^* = 1$, and the unique equilibrium of $\tilde{G}$ is $s_1^* = 0$, even if the distance between u and $\tilde{u}$ is only 2ϵ.

Closeness of Two Games and ϵ-Equilibrium

- However, if u and $\tilde{u}$ are very close, there is a sense in which the equilibria of G are "almost" equilibria of $\tilde{G}$.

Definition

(ϵ-equilibrium) *Given $\epsilon \geq 0$, a mixed strategy $\sigma \in \Sigma$ is called an ϵ-equilibrium if for all $i \in \mathcal{I}$ and $s_i \in S_i$,*

$$u_i(s_i, \sigma_{-i}) \leq u_i(\sigma_i, \sigma_{-i}) + \epsilon.$$

Obviously, when $\epsilon = 0$, an ϵ-equilibrium is a Nash equilibrium in the usual sense.

Continuity Property of ϵ-equilibria

Proposition (1)

Let G be a continuous game. Assume that $\sigma^k \to \sigma$, $\epsilon^k \to \epsilon$, and for each k, σ^k is an ϵ^k-equilibrium of G. Then σ is an ϵ-equilibrium of G.

Proof:

- For all $i \in \mathcal{I}$, and all $s_i \in S_i$, we have

$$u_i(s_i, \sigma^k_{-i}) \leq u_i(\sigma^k) + \epsilon^k,$$

- Taking the limit as $k \to \infty$ in the preceding relation, and using the continuity of the utility functions (together with the convergence of probability distributions under weak topology), we obtain,

$$u_i(s_i, \sigma_{-i}) \leq u_i(\sigma) + \epsilon,$$

  establishing the result.

Closeness of Two Games

- We next define formally the closeness of two strategic form games.

Definition

Let G and G' be two strategic form games with

$$G = \langle \mathcal{I}, (S_i), (u_i) \rangle, \qquad G' = \langle \mathcal{I}, (S_i), (u'_i) \rangle.$$

Then G' is an $\alpha-$approximation to G if for all $i \in \mathcal{I}$ and $s \in S$, we have

$$|u_i(s) - u'_i(s)| \leq \alpha.$$

$\epsilon-$equilibria of Close Games

- The next proposition relates the $\epsilon-$equilibria of close games.

Proposition (2)

If G' is an α-approximation to G and σ is an ϵ-equilibrium of G', then σ is an $(\epsilon + 2\alpha)$-equilibrium of G.

Proof: For all $i \in \mathcal{I}$ and all $s_i \in S_i$, we have

$$\begin{aligned} u_i(s_i, \sigma_{-i}) - u_i(\sigma) &= u_i(s_i, \sigma_{-i}) - u_i'(s_i, \sigma_{-i}) + u_i'(s_i, \sigma_{-i}) - u_i'(\sigma) \\ &\quad + u_i'(\sigma) - u_i(\sigma) \\ &\leq \alpha + \epsilon + \alpha \\ &= \epsilon + 2\alpha. \end{aligned}$$

Approximating a Continuous Game with an Essentially Finite Game

- The next proposition shows that we can approximate a continuous game with an essentially finite game to an arbitrary degree of accuracy.

Proposition (3)

For any continuous game G and any $\alpha > 0$, there exists an "essentially finite" game which is an α-approximation to G.

Proof

- Since S is a compact metric space, the utility functions u_i are uniformly continuous, i.e., for all $\alpha > 0$, there exists some $\epsilon > 0$ such that

$$u_i(s) - u_i(t) \leq \alpha \qquad \text{for all } d(s,t) \leq \epsilon.$$

- Since S_i is a compact metric space, it can be covered with finitely many open balls U_i^j, each with radius less than ϵ (assume without loss of generality that these balls are disjoint and nonempty).
- Choose an $s_i^j \in U_i^j$ for each i, j.
- Define the "essentially finite" game G' with the utility functions u_i' defined as

$$u_i'(s) = u_i(s_1^j, \ldots, s_I^j), \qquad \forall\, s \in U^j = \prod_{k=1}^{I} U_k^j.$$

- Then for all $s \in S$ and all $i \in \mathcal{I}$, we have

$$|u_i'(s) - u_i(s)| \leq \alpha,$$

since $d(s, s^j) \leq \epsilon$ for all j, implying the desired result.

Proof of Glicksberg's Theorem

We now return to the proof of Glicksberg's Theorem. Let $\{\alpha^k\}$ be a scalar sequence with $\alpha^k \downarrow 0$.

- For each α^k, there exists an "essentially finite" α^k-approximation G^k of G by Proposition 3.
- Since G^k is "essentially finite" for each k, it follows using Nash's Theorem that it has a 0-equilibrium, which we denote by σ^k.
- Then, by Proposition 2, σ^k is a $2\alpha^k$-equilibrium of G.
- Since Σ is compact, $\{\sigma^k\}$ has a convergent subsequence. Without loss of generality, we assume that $\sigma^k \to \sigma$.
- Since $2\alpha^k \to 0$, $\sigma^k \to \sigma$, by Proposition 1, it follows that σ is a 0-equilibrium of G.

Discontinuous Games

- There are many games in which the utility functions are not continuous (e.g. price competition models, congestion-competition models).
- We next show that for discontinuous games, under some mild semicontinuity conditions on the utility functions, it is possible to establish the existence of a mixed Nash equilibrium (see [Dasgupta and Maskin 86]).
- The key assumption is to allow discontinuities in the utility function to occur only on a subset of measure zero, in which a player's strategy is "related" to another player's strategy.
- To formalize this notion, we introduce the following set: for any two players i and j, let D be a finite index set and for $d \in D$, let $f_{ij}^d : S_i \rightarrow S_j$ be a bijective and continuous function. Then, for each i, we define

$$S^*(i) = \{s \in S \mid \exists\, j \neq i \text{ such that } s_j = f_{ij}^d(s_i).\} \tag{1}$$

Discontinuous Games

Before stating the theorem, we first introduce some weak continuity conditions.

Definition

Let X be a subset of $\mathbb{R}^n$, X_i be a subset of $\mathbb{R}$, and X_{-i} be a subset of $\mathbb{R}^{n-1}$.

(i) *A function $f : X \to R$ is called upper semicontinuous (respectively, lower semicontinuous) at a vector $x \in X$ if $f(x) \geq \limsup_{k\to\infty} f(x_k)$ (respectively, $f(x) \leq \liminf_{k\to\infty} f(x_k)$) for every sequence $\{x_k\} \subset X$ that converges to x. If f is upper semicontinuous (lower semicontinuous) at every $x \in X$, we say that f is upper semicontinuous (lower semicontinuous).*

(ii) *A function $f : X_i \times X_{-i} \to R$ is called weakly lower semicontinuous in x_i over a subset $X^*_{-i} \subset X_{-i}$, if for all x_i there exists $\lambda \in [0, 1]$ such that, for all $x_{-i} \in X^*_{-i}$,*

$$\lambda \liminf_{x_i' \uparrow x_i} f(x_i', x_{-i}) + (1-\lambda) \liminf_{x_i' \downarrow x_i} f(x_i', x_{-i}) \geq f(x_i, x_{-i}).$$

Discontinuous Games

Theorem (2)

[Dasgupta and Maskin] Let S_i be a closed interval of $\mathbb{R}$. Assume that u_i is continuous except on a subset $S^{**}(i)$ of the set $S^*(i)$ defined in Eq. (1). Assume also that $\sum_{i=1}^n u_i(s)$ is upper semicontinuous and that $u_i(s_i, s_{-i})$ is bounded and weakly lower semicontinuous in s_i over the set $\{s_{-i} \in S_{-i} \mid (s_i, s_{-i}) \in S^{**}(i)\}$. Then the game has a mixed strategy Nash equilibrium.

- The weakly lower semicontinuity condition on the utility functions implies that the function u_i does not jump up when approaching s_i either from below or above.
- Loosely, this ensures that player i can do almost as well with strategies near s_i as with s_i, even if his opponents put weight on the discontinuity points of u_i.

Bertrand Competition with Capacity Constraints

- Consider two firms that charge prices $p_1, p_2 \in [0, 1]$ per unit of the same good.
- Assume that there is unit demand and all customers choose the firm with the lower price.
- If both firms charge the same price, each firm gets half the demand.
- All demand has to be supplied.
- The payoff functions of each firm is the profit they make (we assume for simplicity that cost of supplying the good is equal to 0 for both firms).

Bertrand Competition with Capacity Constraints

- We have shown before that $(p_1, p_2) = (0, 0)$ is the unique pure strategy Nash equilibrium.
- Assume now that each firm has a capacity constraint of 2/3 units of demand:
  - Since all demand has to be supplied, this implies that when $p_1 < p_2$, firm 2 gets $1/3$ units of demand).
- It can be seen in this case that the strategy profile $(p_1, p_2) = (0, 0)$ is no longer a pure strategy Nash equilibrium:
  - Either firm can increase his price and still have $1/3$ units of demand due to the capacity constraint on the other firm, thus making positive profits.
- It can be established using Theorem 2 that there exists a mixed strategy Nash equilibrium.
- Let us next proceed to construct a mixed strategy Nash equilibrium.

Bertrand Competition with Capacity Constraints

- We focus on symmetric Nash equilibria, i.e., both firms use the same mixed strategy.
- We use the cumulative distribution function $F(\cdot)$ to represent the mixed strategy used by either firm.
- It can be seen that the expected payoff of player 1, when he chooses p_1 and firm 2 uses the mixed strategy $F(\cdot)$, is given by

$$u_1(p_1, F(\cdot)) = F(p_1)\frac{p_1}{3} + (1 - F(p_1))\frac{2}{3}p_1.$$

- Using the fact that each action in the support of a mixed strategy must yield the same payoff to a player at the equilibrium, we obtain for all p in the support of $F(\cdot)$,

$$-F(p)\frac{p}{3} + \frac{2}{3}p = k,$$

for some $k \geq 0$. From this we obtain:

$$F(p) = 2 - \frac{3k}{p}.$$

Bertrand Competition with Capacity Constraints

- Note next that the upper support of the mixed strategy must be at $p = 1$, which implies that $F(1) = 1$.
- Combining with the preceding, we obtain

$$F(p) = \begin{cases} 0, & \text{if } 0 \leq p \leq \frac{1}{2}, \\ 2 - \frac{1}{p}, & \text{if } \frac{1}{2} \leq p \leq 1, \\ 1, & \text{if } p \geq 1. \end{cases}$$

Uniqueness of a Pure Strategy Nash Equilibrium in Continuous Games

- We have shown in the previous lecture the following result:
  - Given a strategic form game $\langle \mathcal{I}, (S_i), (u_i) \rangle$, assume that the strategy sets S_i are nonempty, convex, and compact sets, $u_i(s)$ is continuous in s, and $u_i(s_i, s_{-i})$ is quasiconcave in s_i. Then the game $\langle \mathcal{I}, (S_i), (u_i) \rangle$ has a pure strategy Nash equilibrium.
- The next example shows that even under strict convexity assumptions, there may be infinitely many pure strategy Nash equilibria.

Uniqueness of a Pure Strategy Nash Equilibrium

Example

Consider a game with 2 players, $S_i = [0, 1]$ *for* $i = 1, 2$*, and the payoffs*

$$u_1(s_1, s_2) = s_1 s_2 - \frac{s_1^2}{2},$$

$$u_2(s_1, s_2) = s_1 s_2 - \frac{s_2^2}{2}.$$

Note that $u_i(s_1, s_2)$ *is strictly concave in* s_i*. It can be seen in this example that the best response correspondences (which are unique-valued) are given by*

$$B_1(s_2) = s_2, \qquad B_2(s_1) = s_1.$$

Plotting the best response curves shows that any pure strategy profile $(s_1, s_2) = (x, x)$ *for* $x \in [0, 1]$ *is a pure strategy Nash equilibrium.*

Uniqueness of a Pure Strategy Nash Equilibrium

- We will next establish conditions that guarantee that a strategic form game has a unique pure strategy Nash equilibrium, following the classical paper [Rosen 65].

Notation:

- Given a scalar-valued function $f : \mathbb{R}^n \mapsto \mathbb{R}$, we use the notation $\nabla f(x)$ to denote the gradient vector of f at point x, i.e.,

$$\nabla f(x) = \left[\frac{\partial f(x)}{\partial x_1}, \ldots, \frac{\partial f(x)}{\partial x_n}\right]^T.$$

- Given a scalar-valued function $u : \prod_{i=1}^{I} \mathbb{R}^{m_i} \mapsto \mathbb{R}$, we use the notation $\nabla_i u(x)$ to denote the gradient vector of u with respect to x_i at point x, i.e.,

$$\nabla_i u(x) = \left[\frac{\partial u(x)}{\partial x_i^1}, \ldots, \frac{\partial u(x)}{\partial x_i^{m_i}}\right]^T. \quad (2)$$

Optimality Conditions for Nonlinear Optimization Problems

Theorem (3)

(Karush-Kuhn-Tucker conditions) *Let x^* be an optimal solution of the optimization problem*

$$\begin{array}{lll} \text{maximize} & f(x) & \\ \text{subject to} & g_j(x) \geq 0, & j = 1, \ldots, r, \end{array}$$

where the cost function $f : \mathbb{R}^n \mapsto \mathbb{R}$ and the constraint functions $g_j : \mathbb{R}^n \mapsto \mathbb{R}$ are continuously differentiable. Denote the set of active constraints at x^* as $A(x^*) = \{j = 1, \ldots, r \mid g_j(x^*) = 0\}$. Assume that the active constraint gradients, $\nabla g_j(x^*)$, $j \in A(x^*)$, are linearly independent vectors. Then, there exists a nonnegative vector $\lambda^* \in \mathbb{R}^r$ (Lagrange multiplier vector) such that

$$\nabla f(x^*) + \sum_{j=1}^{r} \lambda_j^* \nabla g_j(x^*) = 0,$$

$$\lambda_j^* g_j(x^*) = 0, \qquad \forall\, j = 1, \ldots, r. \tag{3}$$

Optimality Conditions for Nonlinear Optimization Problems

For convex optimization problems (i.e., minimizing a convex function (or maximizing a concave function) over a convex constraint set), we can provide necessary and sufficient conditions for the optimality of a feasible solution:

Theorem (4)

Consider the optimization problem

$$\begin{aligned} \textit{maximize} \quad & f(x) \\ \textit{subject to} \quad & g_j(x) \geq 0, \qquad j = 1, \ldots, r, \end{aligned}$$

where the cost function $f : \mathbb{R}^n \mapsto \mathbb{R}$ and the constraint functions $g_j : \mathbb{R}^n \mapsto \mathbb{R}$ are concave functions. Assume also that there exists some $\bar{x}$ such that $g_j(\bar{x}) > 0$ for all $j = 1, \ldots, r$. Then a vector $x^* \in \mathbb{R}^n$ is an optimal solution of the preceding problem if and only if $g_j(x^*) \geq 0$ for all $j = 1, \ldots, r$, and there exists a nonnegative vector $\lambda^* \in \mathbb{R}^r$ (Lagrange multiplier vector) such that

$$\nabla f(x^*) + \sum_{j=1}^{r} \lambda_j^* \nabla g_j(x^*) = 0,$$

$$\lambda_j^* g_j(x^*) = 0, \qquad \forall\, j = 1, \ldots, r.$$

Uniqueness of a Pure Strategy Nash Equilibrium

- We now return to the analysis of the uniqueness of a pure strategy equilibrium in strategic form games.
- We assume that for player $i \in \mathcal{I}$, the strategy set S_i is given by

$$S_i = \{x_i \in \mathbb{R}^{m_i} \mid h_i(x_i) \geq 0\}, \tag{4}$$

  where $h_i : \mathbb{R}^{m_i} \mapsto \mathbb{R}$ is a concave function.
- Since h_i is concave, it follows that the set S_i is a convex set (exercise!).
- Therefore the set of strategy profiles $S = \prod_{i=1}^{I} S_i \subset \prod_{i=1}^{I} \mathbb{R}^{m_i}$, being a Cartesian product of convex sets, is a convex set.
- Given these strategy sets, a vector $x^* \in \prod_{i=1}^{I} \mathbb{R}^{m_i}$ is a pure strategy Nash equilibrium if and only if for all $i \in \mathcal{I}$, x_i^* is an optimal solution of

$$\begin{aligned} \text{maximize}_{y_i \in \mathbb{R}^{m_i}} \quad & u_i(y_i, x_{-i}^*) \\ \text{subject to} \quad & h_i(y_i) \geq 0. \end{aligned} \tag{5}$$

- We use the notation $\nabla u(x)$ to denote

$$\nabla u(x) = [\nabla_1 u_1(x), \ldots, \nabla_I u_I(x)]^T. \tag{6}$$

Uniqueness of a Pure Strategy Nash Equilibrium

- We introduce the key condition for uniqueness of a pure strategy Nash equilibrium.

Definition

We say that the payoff functions $(u_1, \ldots, u_I)$ are **diagonally strictly concave for** $x \in S$*, if for every* $x^*, \bar{x} \in S$*, we have*

$$(\bar{x} - x^*)^T \nabla u(x^*) + (x^* - \bar{x})^T \nabla u(\bar{x}) > 0.$$

Theorem

Consider a strategic form game $\langle \mathcal{I}, (S_i), (u_i) \rangle$. For all $i \in \mathcal{I}$, assume that the strategy sets S_i are given by Eq. (4), where h_i is a concave function, and there exists some $\tilde{x}_i \in \mathbb{R}^{m_i}$ such that $h_i(\tilde{x}_i) > 0$. Assume also that the payoff functions $(u_1, \ldots, u_I)$ are diagonally strictly concave for $x \in S$. Then the game has a unique pure strategy Nash equilibrium.

Proof

- Assume that there are two distinct pure strategy Nash equilibria.
- Since for each $i \in \mathcal{I}$, both x_i^* and $\bar{x}_i$ must be an optimal solution for an optimization problem of the form (5), Theorem 4 implies the existence of nonnegative vectors $\lambda^* = [\lambda_1^*, \ldots, \lambda_I^*]^T$ and $\bar{\lambda} = [\bar{\lambda}_1, \ldots, \bar{\lambda}_I]^T$ such that for all $i \in \mathcal{I}$, we have

$$\nabla_i u_i(x^*) + \lambda_i^* \nabla h_i(x_i^*) = 0, \tag{7}$$

$$\lambda_i^* h_i(x_i^*) = 0, \tag{8}$$

and

$$\nabla_i u_i(\bar{x}) + \bar{\lambda}_i \nabla h_i(\bar{x}_i) = 0, \tag{9}$$

$$\bar{\lambda}_i h_i(\bar{x}_i) = 0. \tag{10}$$

Proof

- Multiplying Eqs. (7) and (9) by $(\bar{x}_i - x_i^*)^T$ and $(x_i^* - \bar{x}_i)^T$ respectively, and adding over all $i \in \mathcal{I}$, we obtain

$$\begin{aligned} 0 &= (\bar{x} - x^*)^T \nabla u(x^*) + (x^* - \bar{x})^T \nabla u(\bar{x}) \\ &\quad + \sum_{i \in \mathcal{I}} \lambda_i^* \nabla h_i(x_i^*)^T (\bar{x}_i - x_i^*) + \sum_{i \in \mathcal{I}} \bar{\lambda}_i \nabla h_i(\bar{x}_i)^T (x_i^* - \bar{x}_i) \\ &> \sum_{i \in \mathcal{I}} \lambda_i^* \nabla h_i(x_i^*)^T (\bar{x}_i - x_i^*) + \sum_{i \in \mathcal{I}} \bar{\lambda}_i \nabla h_i(\bar{x}_i)^T (x_i^* - \bar{x}_i), \end{aligned} \tag{11}$$

  where to get the strict inequality, we used the assumption that the payoff functions are diagonally strictly concave for $x \in S$.

- Since the h_i are concave functions, we have

$$h_i(x_i^*) + \nabla h_i(x_i^*)^T (\bar{x}_i - x_i^*) \geq h_i(\bar{x}_i).$$

Proof

- Using the preceding together with $\lambda_i^* > 0$, we obtain for all i,

$$\begin{aligned} \lambda_i^* \nabla h_i(x_i^*)^T(\bar{x}_i - x_i^*) &\geq \lambda_i^*(h_i(\bar{x}_i) - h_i(x_i^*)) \\ &= \lambda_i^* h_i(\bar{x}_i) \\ &\geq 0, \end{aligned}$$

  where to get the equality we used Eq. (8), and to get the last inequality, we used the facts $\lambda_i^* > 0$ and $h_i(\bar{x}_i) \geq 0$.

- Similarly, we have

$$\bar{\lambda}_i \nabla h_i(\bar{x}_i)^T(x_i^* - \bar{x}_i) \geq 0.$$

- Combining the preceding two relations with the relation in (11) yields a contradiction, thus concluding the proof.

Sufficient Condition for Diagonal Strict Concavity

- Let $U(x)$ denote the Jacobian of $\nabla u(x)$ [see Eq. (6)]. In particular, if the x_i are all 1-dimensional, then $U(x)$ is given by

$$U(x) = \begin{pmatrix} \frac{\partial^2 u_1(x)}{\partial x_1^2} & \frac{\partial^2 u_1(x)}{\partial x_1 \partial x_2} & \cdots \\ \frac{\partial^2 u_2(x)}{\partial x_2 \partial x_1} & \ddots & \\ \vdots & & \end{pmatrix}.$$

Proposition

For all $i \in \mathcal{I}$, assume that the strategy sets S_i are given by Eq. (4), where h_i is a concave function. Assume that the symmetric matrix $(U(x) + U^T(x))$ is negative definite for all $x \in S$, i.e., for all $x \in S$, we have

$$y^T(U(x) + U^T(x))y < 0, \qquad \forall\, y \neq 0.$$

Then, the payoff functions $(u_1, \ldots, u_I)$ are diagonally strictly concave for $x \in S$.

Proof

- Let x^*, $\bar{x} \in S$. Consider the vector

$$x(\lambda) = \lambda x^* + (1 - \lambda)\bar{x}, \qquad \text{for some } \lambda \in [0, 1].$$

  Since S is a convex set, $x(\lambda) \in S$.

- Because $U(x)$ is the Jacobian of $\nabla u(x)$, we have

$$\begin{aligned} \frac{d}{d\lambda}\nabla u(x(\lambda)) &= U(x(\lambda))\frac{dx(\lambda)}{d(\lambda)} \\ &= U(x(\lambda))(x^* - \bar{x}), \end{aligned}$$

  or

$$\int_0^1 U(x(\lambda))(x^* - \bar{x})d\lambda = \nabla u(x^*) - \nabla u(\bar{x}).$$

Proof

- Multiplying the preceding by $(\bar{x} - x^*)^T$ yields

$$
\begin{aligned}
(\bar{x} - x^*)^T \nabla u(x^*) \quad &+ \quad (x^* - \bar{x})^T \nabla u(\bar{x}) \\
&= \quad -\frac{1}{2} \int_0^1 (x^* - \bar{x})^T [U(x(\lambda)) + U^T(x(\lambda))](x^* - \bar{x}) d\lambda \\
&> \quad 0,
\end{aligned}
$$

where to get the strict inequality we used the assumption that the symmetric matrix $(U(x) + U^T(x))$ is negative definite for all $x \in S$.

6.254 Game Theory with Engineering Applications
Spring 2010