6.254: Game Theory with Engineering Applications February 23, 2010

Lecture 6: Continuous and Discontinuous Games

Lecturer: Asu Ozdaglar

1 Introduction

In this lecture, we will focus on:

- Existence of a mixed strategy Nash equilibrium for continuous games (Glicksberg's theorem).
- Finding mixed strategy Nash equilibria in games with infinite strategy sets.
- Uniqueness of a pure strategy Nash equilibrium for continuous games.

2 Continuous Games

In this section, we consider games in which players may have infinitely many pure strategies. In particular, we want to include the possibility that the pure strategy set of a player may be a bounded interval on the real line, such as [0,1]. We will follow the development of Myerson [3], pp. 140-148.

Definition 1 *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* $u_i : S \mapsto \mathbb{R}$ *are continuous functions.*

A compact metric space is a general mathematical structure for representing infinite sets that can be well approximated by large finite sets. One important fact is that, in a compact metric space, any infinite sequence has a convergent subsequence. Any closed bounded subset of a finite-dimensional Euclidean space is an example of a compact metric space. More specifically, any closed bounded interval of the real line is an example of a compact metric space, where the distance between two points x and y is given by $|x - y|$. In our treatment, we will not need to refer to any examples more complicated than these (see the Appendix for basic definitions of a metric space and convergence notions for probability measures).

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

Theorem 1 (Glicksberg) *Every continuous game has a mixed strategy Nash equilibrium.*

With continuous strategy spaces, the space of mixed strategies Σ is 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.

2.1 Closeness of Two Games and ϵ-Equilibrium

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 $u_i : S \mapsto \mathbb{R}$ and $\tilde{u}_i : S \mapsto \mathbb{R}$ are bounded measurable functions. We may 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.

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 2 (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.

The following result shows that such ϵ-equilibria have a continuity property across games.

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

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

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 [see Eq. (11)], we obtain,

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

establishing the result. $\square$

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

Definition 3 *Let G and G' be two strategic form games with*

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

Assume that the utility functions u_i and u_i' are measurable and bounded. Then G' is an $\alpha-$approximation to G if for all $i \in \mathcal{I}$ and $s \in S$, we have

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

We next relate 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

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

□

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

$$u_i(s) - u_i(t) \leq \alpha,$$

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. □

2.2 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.

3 Discontinuous Games

There are many games in which the utility functions are not continuous (e.g. price competition models, congestion-competition models). For such games existence of a mixed strategy equilibrium can be established under some assumptions by using an existence result by Dasgupta and Maskin [1]-[2]. For more on this topic, please refer to Chapters 12.2 and 12.3 of Fudenberg and Tirole.

In the following examples, we consider games with discontinuous utility functions and find their mixed strategy equilibria.

Example 1 (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).

(a) Find all pure strategy Nash equilibria.

We check all possible candidates to see if there is any profitable unilateral deviation:

- $p_1 = p_2 > 0$: each of the firms has an incentive to reduce their price to capture the whole demand and increase profits.
- $p_1 < p_2$: Firm 1 has an incentive to slightly increase his price.
- $p_1 = p_2 = 0$: Neither firm can increase profits by changing its price unilaterally. Hence $(p_1, p_2) = (0, 0)$ is the unique pure strategy Nash equilibrium.

(b) 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). Show that there does not exist a pure strategy Nash equilibrium. Find a mixed strategy Nash equilibrium.

It is easy to see in this case that the pure 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 Dasgupta-Maskin result that there exists a mixed strategy Nash equilibrium. Let us next proceed to find a mixed strategy Nash equilibrium. 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:

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

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

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

Example 2 (Hotelling Competition): Each of n candidates chooses a position to take on the real line in the interval [0,1]. There is a continuum of citizens, whose favorite positions are uniformly distributed between [0,1]. A candidate attracts votes of citizens whose favorite positions are closer to his position than to the position of any other candidate; if k candidates choose the same position, then each receives the fraction $1/k$ of the votes that the position attracts. The payoff of each candidate is his vote share.

(a) Find all pure strategy Nash equilibria when $n = 2$.

(b) Show that there does not exist a pure strategy Nash equilibrium when $n = 3$.Find a mixed strategy Nash equilibrium.

4 Uniqueness of a Pure Strategy Equilibrium in Continuous Games

We have shown in the previous lecture the following result on the existence of a pure strategy Nash equilibrium: 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.

Then next example shows that even under strict convexity assumptions, there may be infinitely many pure strategy Nash equilibria.

Example 3: Consider a game with 2 players, $S_i = [0, 1]$ for $i = 1, 2$, and the payoffs are given by

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

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

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.

We will next establish conditions that guarantee that a strategic form game has a unique pure strategy Nash equilibrium. We will follow the development of the classical paper by Rosen [4].

We first provide some notation and standard results from nonlinear optimization.

Optimization terminology and background:

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.,

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

(recall our convention that all vectors are column vectors). 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.,

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

We next state necessary conditions for the optimality of a feasible solution of a nonlinear optimization problem. These conditions are referred to as the *Karush-Kuhn-Tucker conditions* in the optimization theory literature.

Theorem 2 (Karush-Kuhn-Tucker conditions) *Let x^* be an optimal solution of 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 continuously differentiable. Denote the set of active constraints at x^* as $A(x^*) = \{j = 1, \dots, r \mid g_j(x^*) = 0\}$. 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

$$\begin{aligned} \nabla f(x^*) + \sum_{j=1}^{r} \lambda_j^* \nabla g_j(x^*) &= 0, \\ \lambda_j^* g_j(x^*) = 0, \qquad &\forall\ j = 1, \dots, r. \end{aligned} \tag{2}$$

You can view the preceding theorem as a generalization of the Lagrange multiplier theorem you might have seen in Calculus for equality constrained optimization problems to inequality constrained optimization problems. The condition that the active constraint gradients, $\nabla g_j(x^*)$, $j \in A(x^*)$, are linearly independent vectors is a regularity condition that eliminates "degenerate" cases where there may not exist Lagrange multipliers. These type of conditions are referred to as *constraint qualifications.* You don't need to worry about these conditions for this class. The condition in Eq. (2) implies that if $\lambda_j^* > 0$, then $g_j(x^*) = 0$, and if $g_j(x^*) > 0$, then $\lambda_j^* = 0$. It is therefore referred to as the *complementary slackness condition* and captures loosely the intuitive idea that a multiplier is used only when the constraint is active and the problem is locally unconstrained with respect to the inactive constraints.

For convex optimization problems (i.e., minimizing a convex function over a convex constraint set, 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 3 *Consider the optimization problem*

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

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 $j = 1, \dots, r$. 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*

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

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

Note that the condition that there exists some $\bar{x}$ such that $g_j(\bar{x}) > 0$ for all $j = 1, \ldots, r$ is a constraint qualification (referred to as the *Slater's constraint qualification*) that is suitable for convex optimization problems.

We now return to the analysis of the uniqueness of a pure strategy equilibrium in strategic form games. In order to discuss the uniqueness of an equilibrium, we provide a more explicit description of the strategy sets of the players. In particular, we assume that for player $i \in \mathcal{I}$, the strategy set S_i is given by

$$S_i = \{x_i \in \mathbb{R}^{m_i} \mid h_i(x_i) \geq 0\}, \tag{3}$$

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 (check as an 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. The following analysis can be extended to the case where S_i is represented by finitely many concave inequality constraints, but we do not do so here for clarity of exposition.

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 the optimization problem

$$\begin{aligned} \text{maximize}_{y_i \in \mathbb{R}^{m_i}} \qquad & u_i(y_i, x_{-i}^*) \\ \text{subject to} \qquad & h_i(y_i) \geq 0. \end{aligned} \tag{4}$$

In the following, we use the notation $\nabla u(x)$ to denote

$$\nabla u(x) = \left[\nabla_1 u_1(x), \ldots, \nabla_I u_I(x)\right]^T, \tag{5}$$

[see Eq. (1)]. We next introduce the key condition for uniqueness of a pure strategy Nash equilibrium.

Definition 4 *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 4 *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. (3), 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 (4), Theorem 3 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

$$\nabla_i u_i(x^*) + \lambda_i^* \nabla h_i(x_i^*) = 0, \tag{6}$$

$$\lambda_i^* h_i(x_i^*) = 0, \tag{7}$$

and

$$\nabla_i u_i(\bar{x}) + \bar{\lambda}_i \nabla h_i(\bar{x}_i) = 0, \tag{8}$$

$$\bar{\lambda}_i h_i(\bar{x}_i) = 0. \tag{9}$$

Multiplying Eqs. (6) and (8) 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

$$\begin{aligned} 0 = (\bar{x} - x^*)^T \nabla u(x^*) \quad &+ \quad (x^* - \bar{x})^T \nabla u(\bar{x}) \qquad (10)\\ &+ \quad \sum_{i \in \mathcal{I}} \lambda_i^* \nabla h_i(x_i^*)^T (\bar{x}_i - x_i^*) + \sum_{i \in \mathcal{I}} \bar{\lambda}_i \nabla h_i(\bar{x}_i)^T (x_i^* - \bar{x}_i) \\ &> \quad \sum_{i \in \mathcal{I}} \lambda_i^* \nabla h_i(x_i^*)^T (\bar{x}_i - x_i^*) + \sum_{i \in \mathcal{I}} \bar{\lambda}_i \nabla h_i(\bar{x}_i)^T (x_i^* - \bar{x}_i), \end{aligned}$$

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

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

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. (7), 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 (10) yields a contradiction, thus concluding the proof. $\square$

Let $U(x)$ denote the Jacobian of $\nabla u(x)$ [see Eq. (5)]. In particular, if the x_i are all 1-dimensional, then $U(x)$ is given by

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

The next proposition provides a sufficient condition for the payoff functions to be diagonally strictly concave.

Proposition 4 *For all $i \in \mathcal{I}$, assume that the strategy sets S_i are given by Eq. (3), 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

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

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

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

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

Recent work has shown that the uniqueness of a pure strategy Nash equilibrium can be established under conditions that are weaker than those given above. This analysis is beyond our scope; for more information on this, see the recent papers [6], [7],[8].

5 Appendix

For more on metric spaces, refer to any textbook on real analysis, for example Rudin [5]. For completeness, we summarize the basic definitions here. A *metric space* is a set M together with a function $d : M \times M \mapsto \mathbb{R}$, that defines the "distance" $d(x,y)$ between any two points x, y in the set. The distance function satisfies the following properties for every x, y, $z \in M$:

- $d(x,y) = d(y,x) \geq 0$,

- $d(x, y) = 0$ if and only if $x = y$,
- $d(x, y) + d(y, z) \geq d(x, z)$.

In a metric space M, a point $y \in M$ is the limit of a sequence of points $\{x^k\}_{k=1}^{\infty} \subset M$ if and only if the distance $d(x^k, y) \to 0$ as $k \to \infty$. An *open ball* of radius ϵ around a point x, denoted by $B(x, \epsilon)$, is the set of all points in the metric space that have a distance less than ϵ from x, i.e.,

$$B(x, \epsilon) = \{y \mid d(x, y) < \epsilon\}.$$

A set S is an *open* subset of a metric space M if and only if for every $x \in S$, there exists some $\epsilon > 0$ such that $B(x, \epsilon) \subset S$. A metric space is compact if and only if every collection of open sets that "covers" M (i.e., their union includes all of M), has a finite subcollection that also covers M.

When there are infinitely many actions in the set S_i, a mixed strategy for player i can no longer be described by just listing the probability of each individual action (i.e., by a finite dimensional probability vector). For example, suppose that S_i is the interval $[0, 1]$. If player i selected his action from a uniform probability distribution over the interval [0,1], then each individual action in [0,1] would have zero probability; but the same would be true if he selected his action from a uniform probability distribution over the interval [0.5,1]. To describe a probability distribution over S_i, we must list the probabilities of subsets of S_i. Unfortunately, for technical reasons, it may be mathematically impossible to consistently assign probabilities to all subsets of an infinite set, so some weak restriction is needed on the class of subsets whose probabilities can be meaningfully defined. These are called the *measurable sets*. Here, we let the measurable subsets of S and of each set S_i be the smallest class of subsets that includes all open subsets, all closed subsets, and all finite or countably infinite unions and intersections of sets in the class. These are the *Borel subsets* (and they include essentially all subsets that could be defined without the use of very sophisticated mathematics). Let $\mathcal{B}_i$ denote the set of such measurable or Borel subsets of S_i.

Let Σ_i denote the set of probability distributions over S_i, i.e., $\sigma_i \in \Sigma_i$ if and only if σ_i is a function that assigns a nonnegative number $\sigma_i(Q)$ to each Q that is a Borel subset of S_i,

$\sigma(S_i) = 1$, and for any countable collection $(Q^k)_{k=1}^{\infty}$ of pairwise-disjoint Borel subsets of S_i,

$$\sigma_i\left(\bigcup_{k=1}^{\infty} Q^k\right) = \sum_{k=1}^{\infty} \sigma_i(Q^k).$$

We define convergence for mixed strategies by assigning the *weak topology* to Σ_i. Two implications of this topology are important for our purposes:

- The set of probability distributions Σ_i is a compact metric space, i.e., every sequence $\{\sigma_i^k\} \subset \Sigma_i$ has a convergent subsequence.

- A sequence $\{\sigma_i^k\} \subset \Sigma_i$ of mixed strategies converges to $\sigma_i \in \Sigma_i$ if and only if for all continuous functions $f : S_i \mapsto \mathbb{R}$, we have

$$\lim_{k\to\infty} \int_{S_i} f(s_i) d\sigma_i^k(s_i) = \int_{S_i} f(s_i) d\sigma_i(s_i). \tag{11}$$

The preceding condition asserts that if $\tilde{s}_i^k$ is a random strategy drawn from S_i according to the σ_i^k distribution, and $\tilde{s}_i$ is a random strategy drawn from S_i according to the σ_i distribution, then the expected value of $f(\tilde{s}_i^k)$ must converge to the expected value of $f(\tilde{s}_i)$ as $k \to \infty$.

A function $g : S \mapsto \mathbb{R}$ is called *(Borel) measurable* if for every scalar t, the set $\{x \in S \mid g(x) \geq t\}$ is a Borel subset of S. A function $g : S \mapsto \mathbb{R}$ is bounded if there exists some scalar K such that $|g(x)| \leq K$ for all $x \in C$. To be able to define expected utilities, we must require that player's utility functions are measurable and bounded in this sense. Note that this assumption is weaker than the continuity of the utility functions. With this assumption, the utility functions u_i are extended from $S = \prod_{j=1}^{I} S_j$ to the space of probability distributions $\Sigma = \prod_{j=1}^{I} \Sigma_j$ as follows:

$$u_i(\sigma) = \int_C u_i(s) d\sigma(s),$$

where $\sigma \in \Sigma$.

References

[1] Dasgupta, P. and Maskin E., "The existence of equilibrium in discontinuous economic games. 1: Theory," *Review of Economic Studies*, vol. 53, pp. 1-26, 1986.

[2] Dasgupta, P. and Maskin E., "The existence of equilibrium in discontinuous economic games. 2: Applications," *Review of Economic Studies*, vol. 53, pp. 27-42, 1986.

[3] Myerson R., Game Theory: Analysis of Conflict, Harvard University Press, Cambridge, MA, 1991.

[4] Rosen J. B., "Existence and uniqueness of equilibrium points for concave N-person games," *Econometrica*, vol. 33, no. 3, 1965.

[5] Rudin W., Principles of Mathematical Analysis, McGraw-Hill, 1976.

[6] Simsek, A., Ozdaglar, A., and Acemoglu, D., "A generalized Poincare-Hopf Theorem for compact nonsmooth regions," *Mathematics of Operations Research*, vol. 32, no. 1, pp. 193-214, 2007.

[7] Simsek, A., Ozdaglar, A., and Acemoglu, D., "Local indices for degenerate variational inequalities, to appear in *Mathematics of Operations Research*, 2008.

[8] Simsek, A., Ozdaglar, A., and Acemoglu, D., "Uniqueness of generalized equilibrium for box constrained problems and applications," *Proc. of Allerton Conference*, September 2005.