FDPE Microeconomic Theory II

Spring 2007

by

Hannu Vartiainen

1 Introduction

- General Equilibrium theory: competition is good.
- No real interaction between the decision makers. Everyone takes the external parameters - such as prices - as given, and optimizes without paying attention on (i) how his behavior affects the parameters (in the equilibrium it does not), (ii) where do the parameters come from.
- General equilibrium theory does not have anything to about how decision makers *interact*.
- Is interaction relevant?

- *Game theory* is a collection of analytical tools designed to help us understand the phenomena we observe when decision makers interact.
- The central assumptions that underlie the theory are that decision makers pursue well-defined exogenous objects (they are *rational*), and take into account that others do too (think *strategically*).
- A *game* describes the environment where the interaction takes place. The concept of game is self-containing: all the relevant information is spelled out, i.e. the rules of interaction, who interact, what they know, and their preferences.
- Interaction of rational decision makers within a game is captured by the concept of *equilibrium*.

• An equilibrium is a steady state where all possible adjustment processes have already taken place. An equilibrium concept answers the question of what constraints does decision makers' rationality, their knowledge of other players rationality and so on, impose on their collective behavior.

- Game theory is a *language*. The models of game theory do not aim to *describe* the reality but they can be useful for us to *think* about it.
- Applications should be viewed as thought experiments, aimed at straightening the conceptual problems in our heads.
- Applications abound:
 - Competition.
 - Politics.
 - Bargaining.
 - Organizations.

- Psychology.
- etc...
- Often regarded as "mathematical". However, formality only aims at making the language clear, crisp, and transparent.
- A high level of abstraction of the model allows it to be applied in many contexts.
- The aim of this course: To develop the language of game theory, and demonstrate how it can be applied when thinking applications.

2 Rationality

- Decision maker is rational in the sense that he fully understands the choice problem he is facing:
 - A set of actions A that the decision maker can take.
 - An exhaustive set of possible consequences C.
 - A consequence function $g: A \to C$ specifying which actions lead to which consequences.
 - Complete, transitive preferences \succeq on the set C.
- There is a real valued *utility function* u on A that represents preferences, and can be used instead of the preferences. That is, $u(g(a)) \ge u(g(b))$

if and only if $g(a) \succeq g(b)$ for all $a, b \in A$. In applications we use utility functions. However, it is important that the preferences are the true primitive (why?).

- Rationality implies that the decision maker maximizes his payoffs; if u(g(b)) > u(g(a)), for some a, b ∈ A, then b is not chosen.
- To model decision making under *uncertainty*, we assume that decision maker's preferences also satisfy von Neumann -Morgenstern axioms. That is, if the consequence does not only depend on the decision maker's choice but also on a stochastic component, then the decision maker is assumed to behave *as if* he maximizes his expected utility, for some utility function *u*.

- More formally:
 - Let the consequence depend not only decision maker's action but also a state in some finite set Ω .
 - If the decision maker takes action a, and state $\omega \in \Omega$ materializes, then the decision maker's payoff is u(g(a, w)).
 - If the uncertainty is captured by a probability distribution p on Ω , the decision maker maximizes his expected payoff

 $\sum_{\omega\in\Omega}p(\omega)u(g(a,\omega)).$

3 Strategic games

- A strategic game is a model of interactive decision making when the decision makers, "players", choose once and for all and simultaneously their actions. The actions need not be physically simultaneous, only independent in players' minds.
- The strategic game is a tuple $\langle I, (u_i)_{i \in I}, (A_i)_{i \in I} \rangle$ where
 - I is a finite set of players i, j, \ldots .
 - A_i is the set of player *i*'s possible actions.
 - $u_i : \times_{i \in I} A_i \to \mathbb{R}$ is player *i*'s von Neumann-Morgenstern utility function.

- That is, any combination a = (a_i)_{i∈I} of actions of players generates utility u_i(a) for any player i. Rationality now means that each player i maximizes his payoff by choosing his action in A_i optimally.
- We assume that rationality is *common knowledge:* All players are rational, all players know that all players are rational, all players know that all players know that
- The requirement that player j's preferences are now defined over A = ×_{i∈I}A_i rather than A_i distinguishes a strategic form game from a single player decision problem. Or from the Walrasian markets: Each player (may) not only care about his own action but also the actions of all other players.

• Concept check: By von Neumann -Morgenstern utilities we mean that the underlying preferences over *lotteries* on $A = \times_{i \in I} A_i$ meet the von Neumann - Morgenstern axioms. If they do, then each player *i* acts *as if* there is a payoff function u_i with respect to which he maximizes his *expected* payoff.

3.1 Dominated strategies

• What does rationality alone imply for behavior in strategic form games?

Prisoner's Dilemma: Let $I = \{1, 2\}, A_i = \{C, D\}$ for all i, and let payoffs be determined as follows:

Whatever strategy the other player chooses, it is player i's best response to choose D and not C. Thus C is a dominated strategy.

• A rational player would never choose an action that he could not justify, i.e. that is not a best response to *any* of the strategies of the other players.

Definition (strict dominance): A strategy a_i is *strictly dominated* if there is another strategy a'_i such that

$$u_i(a'_i, a_{-i}) > u_i(a_i, a_{-i}), \text{ for all } a_{-i} \in \times_{j \neq i} A_j.$$

- If strategy a_i is strictly dominated for i, then i would never lose by choosing the strategy that dominates it. Thus rationality of i *implies* that such strategy would never be chosen.
- This argument can be pushed further. Since rationality is common knowledge, it is known that no player uses dominated strategies. Since this fact is known, players face a new decision problem, in which new strategies become eliminated by a second use of rationality. The elimination process can be continued until no further strategies can be eliminated. The survived strategies are called *iteratively undominated strategies*.

Definition (iterative dominance): Set $X \subseteq A$ of strategies are *iteratively* undominated if $X_i = A_i \setminus \bigcup_{k=1}^{\infty} D_i^k$ where $D_i^k \subset A_i, k = 1, ...,$ is defined such that $a_i \in D_i^k$ if and only if

$$u_i(a'_i, a_{-i}) > u_i(a_i, a_{-i}), \text{ for all } a_{-i} \in \times_{j \neq i} A_j \setminus \cup_{l=1}^{k-1} D_j^l,$$

nd $D^0 = \emptyset$.

а

- Strict dominance is attractive since it is directly *implied by rationality*: Common knowledge of rationality means that players would only use strategies in X.
- In two-player games, the converse is also true. Rational players may use any stragies in X.

Example (Cournot competetion): Let us model the two-firm Cournot model as a game $\langle \{1, 2\}, (u_i), (A_i) \rangle$, where $A_i = \mathbb{R}_+$ and, for any $(a_1, a_2) \in A_1 \times A_2$,

$$u_1(a_1, a_2) = \beta(a_1, a_2)a_1,$$

 $u_2(a_1, a_2) = \beta(a_1, a_2)a_2.$

where β is a linear (inverse) demand function

$$\beta(a_1, a_2) = 1 - (a_1 + a_2).$$

Thus i's payoff is

$$u_1(a_1, a_2) = (1 + a_2)a_1 - a_1^2.$$

Taking the derivative gives the effect of a marginal increase in a_i on i's payoff:

$$\frac{\partial u_i(a_i, a_j)}{\partial a_i} = 1 - a_j - 2a_i. \tag{1}$$

If (1) is positive (negative) under (a_i, a_j) , then marginally increasing (decreasing) a_i increases *i*'s payoff. If this holds continuously in the interval [a, b] of *i*'s choices under a_j , then increasing a_i from *a* to *b* increases *i*'s payoff.

• By (1), $a_i = 1/2$ strictly dominates any $a_i > 1/2$, given that $a_j \ge 0$. Thus

$$D_i^1 = \left\{ a_i : a_i > \frac{1}{2} \right\}$$

Applying the symmetric argument to j,

$$A_j \setminus D_j^1 = \left\{ a_j : \mathbf{0} \le a_j \le \frac{1}{2} \right\}.$$

By (1), $a_i = 1/2 - (1/2)^2$ strictly dominates any $a_i < 1/2 - (1/2)^2$,

given that $0 \le a_j \le 1/2$. Thus

$$D_i^2 = \left\{ a_i : a_i < \frac{1}{2} - \left(\frac{1}{2}\right)^2 \right\}.$$

Applying the symmetric argument to j,

$$A_j \setminus \left(D_j^1 \cup D_j^2 \right) = \left\{ a_j : \frac{1}{2} - \left(\frac{1}{2} \right)^2 \le a_j \le \frac{1}{2} \right\}.$$

By (1), $a_i = 1/2 - (1/2)^2 + (1/2)^3$ strictly dominates any $a_i > 1/2 - (1/2)^2 + (1/2)^3$, given that $1/2 - (1/2)^2 \le a_j \le 1/2$. Thus

$$D_i^3 = \left\{ a_i : a_i > \frac{1}{2} - \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 \right\}.$$

Applying the symmetric argument to j,

$$A_j \setminus \left(D_j^1 \cup D_j^2 \cup D_j^3 \right) = \left\{ a_j : \frac{1}{2} - \left(\frac{1}{2}\right)^2 \le a_j \le \frac{1}{2} - \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 \right\}.$$

Continuing this way for $k \pmod{3}$ steps, we get

$$A_{j} \setminus \left(\cup_{l=1}^{k} D_{j}^{l} \right) = \begin{cases} \frac{1}{2} - \left(\frac{1}{2}\right)^{2} + \left(\frac{1}{2}\right)^{3} - \dots - \left(\frac{1}{2}\right)^{k-1} \\ a_{j} : & \leq a_{j} \leq \\ \frac{1}{2} - \left(\frac{1}{2}\right)^{2} + \left(\frac{1}{2}\right)^{3} - \dots + \left(\frac{1}{2}\right)^{k} \end{cases}$$

Letting k go to infinity, both the end points of the interval converge to

$$\frac{1/2}{1-(1/2)^2} - \frac{(1/2)^2}{1-(1/2)^2} = \frac{1}{3}.$$

Thus

$$\left(\frac{1}{3},\frac{1}{3}\right)$$

is the unique strategy pair that survives the iterated elimination of strictly dominated strategies.

- However, this strict dominance turns out to be a weak restriction on strategies. In many interesting games, it does not imply anything.
- A slightly stronger restriction is weak dominance:

Definition (weak dominance): A strategy a_i is *weakly dominated* if there is another strategy a'_i such that

$$u_i(a'_i, a_{-i}) \ge u_i(a_i, a_{-i}), \text{ for all } a_{-i} \in \times_{j \ne i} A_j,$$

with at least one strict inequality.

• It is always at least as profitable for i to choose a'_i than a_i , and sometimes more profitable, no matter what strategy the other players choose.

- The motivation for weak dominance is insurance: A risk sensitive player would never choose a weakly dominated actions since he can always do at least as well by doing something else.
- Weak dominance can also be iterated. However, it may not always be obvious how to do it.

Example 1: Consider the game



Step 1: Delete a_2

	b_1	b_2	b_3	b_{4}
a_1	2,1	2,1	-1, 0	-1, 0
a_3	3,1	0,0	3, 1	0,0
a_{4}	0,0	1,3	0,0	1,3

Step 2: Delete b_3 and b_4

Step 3: Delete a_4

Step 4: Delete b_2

Step 5: Delete a_1 . The unique remaining strategy is (a_3, b_1) .

- The key problem with weak dominance is that the order of iteration may affect the conclusion.
- Iteratively undominated strategies always exist (why?). However, dominance typically has only limited power.

Example (Battle of Sexes):



3.2 Nash equilibrium

 Since iterative dominance does not have much power, rationality alone does not provide much insight into reasonable individual behavior in social situations. To constraint further individuals' behvior, something stronger needs to be assumed.

- Nash equilibrium is a more stringent condition on strategic behavior. It is not derived through recursive reasoning, which is the case with dominated strategies, but through counterfactual one. Players ask not only what are their best responses against the *potential* best responses of the other players but also what are their best responses if the best responses of the other other players are *correct*.
- Nash equilibrium requires that each player's strategy be a *best response* to the other players' strategies. In the battle of sexes -game, there are two distinct Nash equilibria.
- However, there need not be determined choices that satisfy the criteria for the Nash equilibrium.

• In the matching pennies -game, one of the players always wants to change his strategy; there is no rest point in the associated *best response* - *dynamics*.

Example (Matching Pennies):

3.2.1 Mixed strategies

• To avoid the cyclical best response dynamics, we need to enlarge the class of strategic alternatives available to the players.

- We allow players to randomize over their choices. A randomized strategy is called *mixed*. A mixed strategy is a lottery over the choices of a player.
- The set of lotteries on A_i is denoted by Δ_i , with a typical element σ_i . Denote by σ a typical element of $\times_i \Delta_i$.
- In mixed strategy σ_i , player *i* assigns probability $\sigma_i(a_i)$ to choice a_i . If σ_i is degenerate, i.e. $\sigma_i(a_i) = 1$ for some a_i , then σ_i is *pure*.
- Player *i*'s expected payoff, when all players adhere to mixed strategy σ , can now be written (assuming finite A),

$$u_i(\sigma) = \sum_{a \in A} \prod_{j \in I} \sigma_j(a_j) u(a).$$

• Nash equilibrium (by Nash, 1951) is the central solution concept in game theory.

Nash Equilibrium: Mixed strategy σ^* forms a Nash equilibrium in game $\langle I, (u_i), (A_i) \rangle$ if

$$u_i(\sigma_i^*, \sigma_{-i}^*) \ge u_i(\sigma_i, \sigma_{-i}^*), \text{ for all } \sigma_i \in \Delta_i, \text{ for all } i \in I.$$

• If a single strategy survives iterative elimination of strictly dominated strategies, then this strategy must form a *unique* Nash equilibrium.

Example (Cournot comp., cont.): Recall condition (1), the derivative of firm *i*'s payoff given the other firm's choice in the Cournot competition model.

Firm i's payoff is maximized when

$$\frac{\partial u_i(a_i, a_j)}{\partial a_i} = 1 - a_j - 2a_i = 0.$$

From this first order condition we derive the two best response functions:

$$BR_{1}(a_{2}) = \frac{1}{2} - \frac{a_{2}}{2}, \qquad (2)$$
$$BR_{2}(a_{1}) = \frac{1}{2} - \frac{a_{1}}{2}.$$

The Nash equilibrium (a_1^*, a_2^*) satisfies.

$$a_1^* = BR_1(a_2^*),$$

 $a_2^* = BR_2(a_1^*).$

By equation (2) it follows that

$$a_1^* = \frac{1}{2} - \frac{a_2^*}{2},$$

 $a_2^* = \frac{1}{2} - \frac{a_1^*}{2}.$

Thus we obtain

$$(a_1^*, a_2^*) = \left(\frac{1}{3}, \frac{1}{3}\right).$$

The equilibrium coincides with the unique strategy pair that survives the itarated elimination of dominated strategies. Is this a coincidence?

Equilibrium payoffs are

$$u_1\left(\frac{1}{3},\frac{1}{3}\right) = u_2\left(\frac{1}{3},\frac{1}{3}\right) = \left(1-\frac{2}{3}\right)\frac{1}{3} = \frac{1}{9}.$$

• A Nash equilibrium, and only a Nash equilibrium has the property that

rational players predict correctly what the other players choose, predicts that their rational opponents predict it, and so on.

- On the one hand, playing Nash equilibrium can be agreed upon, since no-one can benefit by unilaterally deviating from the agreement. On the other hand, any viable agreement must induce Nash equilibrium play since otherwise someone wants to renege on the agreement. Thus players can agree on playing certain strategy if and only if this strategy forms a Nash equilibrium.
- Pure Nash equilibrium puts positive probability only to a single strategy. Such equilibria commonly fail to exists (cf. mathcing pennies).
- The existence of Nash equilibrium in mixed strategies is the key result in game theory. It is based on a fixed point theorem.

- Kakutani's Fixed Point Theorem: Let X be a compact, convex subset of a finitely dimensional Euclidean space. Let $F : X \to X$ be a set valued function for which (i) F(x) is non-empty and convex for all $x \in X$, (ii) the graph of F is closed. Then there is a fixed point $x^* \in X$ such that $x^* \in F(x^*)$.
 - Function F has a closed graph if any sequence {x^k} converging to y supports a sequence {f(x^k)}_k such that f(x^k) ∈ F(x^k) for all k = 0,..., that converges in F(y).
 - Sketch of the proof the existence of Nash equilibrium: Denote the *best response function* of *i* by

$$BR_i(\sigma_{-i}) = \{\sigma_i \in \Delta_i : u_i(\sigma_i, \sigma_{-i}) \ge u_i(\sigma'_i, \sigma_{-i}), \text{ for all } \sigma'_i \in \Delta_i\}.$$

- Let BR(σ) = ×_iBR_i(σ_{-i}) for all σ ∈ ×_iΔ_i. Now BR is a set valued functions from ×_iΔ_i to ×_iΔ_i.
- We argue that BR meets the conditions imposed in Kakutani if A is finite. It is easy verify that $BR_i(\sigma_{-i})$ is nonempty and convex.
- To see that BR has a closed graph, it suffices that the graph is closed in each of its dimension. Suppose that there is a sequence $\{\sigma_{-i}^k\}$ of strategies, converging to σ_{-i} such that $\{\sigma_i^k\}$, where $\sigma_i^k \in BR_i(\sigma_{-i}^k)$ for all k, would converge to $\sigma_i \notin BR_i(\sigma_{-i})$. Then there is σ'_i and $\varepsilon > 0$ such that

$$u_i(\sigma'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) > \varepsilon.$$

Since u_i is continuous in probabilities, there is k' such that for all k > k',

$$u_i(\sigma'_i, \sigma_{-i}) - u_i(\sigma^k_i, \sigma^k_{-i}) > \varepsilon/2.$$

Similarly, there is k'' > k' such that for all k > k'',

$$u_i(\sigma'_i, \sigma^k_{-i}) - u_i(\sigma^k_i, \sigma^k_{-i}) > \varepsilon/3.$$

But this contradicts $\sigma_i^k \in BR_i(\sigma_{-i}^k)$.

- Since each Δ_i is a finitely dimensional standard simplex, ×_iΔ_i is a convex set.
- Thus, by Kakutani's Theorem, there is σ^* such that $\sigma^* \in BR(\sigma^*)$, i.e. $\sigma_i^* \in BR_i(\sigma_{-i}^*)$ for all *i*.

- **Theorem (Nash 1950):** If A is finite, then the game $\langle I, (u_i), (A_i) \rangle$ has a Nash equilibrium.
 - Do people randomize? Mixed strategies have alternative interpretations: steady state of a learning process, beliefs that derived from a distribution,...

.

- Finiteness of A is needed. Consider the game $I = \{1, 2\}, A_i = \mathbb{Z}$, and $u_i(a_1, a) = 1$ if $a_i > a_2$, and = 0 otherwise.
- A very useful feature of mixed strategy equilibria is the following: σ^* is a Nash equilibrium only if every pure action in the support of σ_i , i.e. in

 $\{a_i : \sigma_i^*(a_i) > 0\}$, is a best response to σ_{-i}^* and hence yields *i* the same payoff. That is, if $\sigma_i^*(a_i) > 0$, for some $a_i \in A_i$, then

$$u_i(a_i, \sigma^*_{-i}) \ge u_i(a'_i, \sigma^*_{-i}), \text{ for all } a'_i \in A_i.$$

• This follows from the linearity of *u* in probabilities. For if the claim would not hold, *i* could improve his payoff by moving probability mass to those actions that generate him a higher payoff (keeping other players strategies fixed).

Example (Matching Pennies, cont.): The unique mixed strategy Nash equilibrioim both players mix 50/50 between H and T.

Example (Advertizing): Two firms compete over consumer's attention through (noninformative) advertizment campaings. The firm who invests more to

its campaign sells the good to the consumer. Campaigns are chosen simultaneously. The consumer pays $1000 \in$ for the good. If firm *i* wins its payoff is $1000 - a_i$, and if it loses it is $-a_i$, where *i*'s investment a_i can be any nonnegative number. In the case of a tie, the consumer randomizes 50/50 from who to buy.

There are no pure Nash equilibria. For if there is, then either (i) $a_1 = a_2$, or (ii) $a_i > a_j$ for some choice of i, j. In case (i), i would deviate by increasing a_i slightly above a_j . In case (ii), i would benefit from slightly decreasing a_i so that it is still above a_j .

In a continuous, symmetric equilibrium (it can be shown to be unique), firm *i* chooses an atomless mixed strategy over some interval [a, b]. Since the strategy is mixed, the firm must be indifferent between any choice in the intervall. Denoting the mixed strategy of *i* by a c.d.f. Σ_i with density σ , the probability of *i* winning when choosing action a_i is the probability of j choosing less than a_i , i.e. $\Sigma_j(a_i)$. Since $\Sigma_j(\cdot)$ is an atomless c.d.f., $\Sigma_j(a) = 0$, and thus the probability of winning at a is zero. This implies that a cannot be above 0 since otherwise i would rather choose $a_i = 0$ than $a_i = a$, which is against our assumption. Since choosing $a_i = a$ is part of i's Nash equilibrioum strategy, and he is indifferent between all his choices, his expected payoff must be zero. Since by choosing b firm i guarantee positive payoff if b is below 1000, and i is indifferent in all choices [a, b], is must be that b = 1000. Thus in Nash equilibrium,

 $\Sigma_j(a_i) \cdot 1000 - a_i = 0$, for all $a_i \in [0, 1000]$.

That is $\sigma_j(a_i) = 1/1000$, for all $a_i \in [0, 1000]$, for all i, j.

The expected aggregate costs of advertizing are

$$\int_{0}^{1000} a\sigma_{1}(a)da + \int_{0}^{1000} a\sigma_{2}(a)da = \frac{2}{1000} \int_{0}^{1000} ada = \frac{2}{1000} \cdot \frac{1000^{2}}{2} = \frac{2}{1000} \cdot \frac{1000^{2}}{2}$$

Thus the expected costs of advertizing match exactly with the consumer surplus, destroing any social surplus.

- The above argument can be applied to lobbying, doping, in sports, cosmetic surgery,... .
- Things become more problematic if there are many equilibria. Then little can be said which will be selected. There is no obvious process that leads to the Nash equilibrium.

3.2.2 Multiplicity of Equilibria

Example (Coordination game): Equilibria may be Pareto rankable.

It would be natural to assume that Pareto dominant equilibrium be played, especially if players could coordinate before they play. But this is not obvious. As argued by Aumann (1990), pre-play communication can only have value if it conveys new information on how players will behave in the game. In the Coordination game, a player always wants his opponent to choose D strategy regardless of what they themselves intend to choose. Hence messages convey no information about what players intend to do,

only about what they want their opponents to do. This information is already common knowledge to the players, and so communication cannot affect the outcome of the game.

• A game may exhibit mixed strategy Nash equilibria together with pure strategy equilibria.

Example (Burning money) Consider a coordination game

This game has two pure Nash equiilibria (U, L) and (D, R). Player 1 prefers the first one, 2 the latter one. There is no way to choose among them.

Allow now player 1 to burn one unit of money before playing the coordination game. However, the strategies cannot be implemented personally but through a mediator. Both players submit simultaneously complete instructions of how to choose to a mediator who, after seeing both instructions, implementes them.

A strategy of 1 specifies whether to burn the money or not, and whether to choose U or D.

A strategy of 2 specifies, conditional on 1 burning the money, whether to choose L or R .

Thus both players have four options (in the matrix below, L, R means 2

chooses L if 1 burns the money, and R if not, and so on).

		2			
		L, L	L,R	R,L	R,R
	$Burn, \ U$	2,1	2,1	-1, 0	-1,0
1	$Burn, \ D$	-1, 0	-1, 0	0,3	0,3
	Not burn, U	3,1	0,0	3,1	0,0
	Not burn, D	0,0	1,3	0,0	1,3

Compare this to Example 1!

Example (Hawk-Dove): Each player wants to be a hawk against a dove, but not hawk against a hawk.

3.2.3 Zero-sum games

- A two player game ({1,2}, (u_i), (A_i)) has a zero-sum property if u₁(a₁, a₂) = -u₂(a₁, a₂), for all (a₁, a₂) ∈ A₁ × A₂. This implies, in particular, that u₁(σ₁, σ₂) = -u₂(σ₁, σ₂), for all σ ∈ Δ₁ × Δ₂.
- Many games have the zero-sum property, sports, contests, politics,... .
- Player 1's *minimax* payoff is

$$ar{u}_1 = \min_{\sigma_2 \in \Delta_2} \max_{\sigma_1 \in \Delta_1} u_1(\sigma_1, \sigma_2),$$

and player 2's minimax payoff is

$$\overline{u}_2 = \min_{\sigma_1 \in \Delta_1} \max_{\sigma_2 \in \Delta_2} u_2(\sigma_1, \sigma_2).$$

• The next proposition shows that in zero-sum game every Nash equilibrium generates players the same payoffs. Thus the payoffs of the game are uniquely specified.

Proposition: Every Nash equilibrium induces payoffs (\bar{u}_1, \bar{u}_2) .

- Suppose that (σ_1^*, σ_2^*) is a Nash equilibrium of the game. Then $u_2(\sigma_1^*, \sigma_2^*) \ge u_2(\sigma_1^*, \sigma_2)$, for all σ_2 . Since $u_1 = -u_2$, we have $u_1(\sigma_1^*, \sigma_2^*) \le u_1(\sigma_1^*, \sigma_2)$, for all σ_2 . Thus $u_1(\sigma_1^*, \sigma_2^*) \le \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \le \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$. Further, $u_1(\sigma_1^*, \sigma_2^*) \ge u_1(\sigma_1, \sigma_2^*)$ for all σ_1 . In particular, $u_1(\sigma_1^*, \sigma_2^*) \ge \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$, for all σ_1 , so that $u_1(\sigma_1^*, \sigma_2^*) \ge \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$. Thus $u_1(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2) = \overline{u_1}$. Similarly for player 2.
- Note that since $\bar{u}_1 = -\bar{u}_2 = -\max_{\sigma_2} \min_{\sigma_1} u_2(\sigma_1, \sigma_2) = \min_{\sigma_2} \max_{\sigma_1} -u_2(\sigma_1, \sigma_2) = \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$, we have $\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$.

• Consider the game of tennis. Let p_{fb} denote the probability that the server (S) wins a point when he serves to forehand (f) when the responder (R) anticipates a serve to the backhand (b) corner, and so on. Assume that $0 \le p_{ff}, p_{bb} < p_{fb}, p_{bf} \le 1$. That is, that the server wins with a higher probability if he manages to serve to the corner to which the responder does not expect the ball to come. The payoffs are

		R	
		Fore	Back
S	Fore	$p_{ff}, 1 - p_{ff}$	$p_{fb}, 1 - p_{fb}$
	Back	$p_{bf}, 1 - p_{bf}$	$p_{bb}, 1 - p_{bb}$

• The probability of winning when S serves to forehand is $\sigma_R(Fore)p_{ff} + \sigma_R(Back)p_{fb}$ and to backhand $\sigma_R(Fore)p_{bf} + \sigma_R(Back)p_{bb}$. In the unique mixed strategy Nash equilibrium, player S randomizes:

$$\sigma_R(Fore)p_{ff} + \sigma_R(Back)p_{fb} = \sigma_R(Fore)p_{bf} + \sigma_R(Back)p_{bb}.$$

Thus, in equilibrium, the probability of winning a point from either corners is equal, for *any* winning probability system p_{ff} , p_{bb} , p_{fb} , p_{bf} . A testable hypothesis!

3.3 Correlated equilibrium

- In Nash equilibrium, players actions are independent.
- One of the motivations of mixed strategy is that pure actions are actually conditional on random, independent signals.
- What if the signals are not independent?

Correlated equilibrium: Probability distribution π over A is a correlated equilibrium of strategic game $\langle I, (u_i)_{i \in I}, (A_i)_{i \in I} \rangle$ if

$$\sum_{a_{-i}} \pi(a_i, a_{-i}) u_i(a_i, a_{-i}) \geq \sum_{a_{-i}} \pi(a_i, a_{-i}) u_i(a'_i, a_{-i}),$$

for all i , for all a_i , for all a'_i .

- That is, player i should not be able to gain by disobeying the recommendation to play a_i if every other played obeys the recommendation.
- Every Nash equilibrium is a correlated equilibrium but not vice versa. The Battle of Sexes game indeces a correlated equilibrium where plauers get expected payoff (2/3, 2/3) (contrast this to the mixed Nash equilibrium). This is achieved by tossing a fair coin, an conditioning the pure equilibria to he result.

• Correlated equilibria have nice geometric features.

Theorem: Let A be finite. The set of correlated equilibria is convex.

• Rewriting the equilibrium condition

$$\sum_{a_{-i}} \pi(a_i, a_{-i}) \delta(i, a_i, a'_i, a_{-i}) \ge 0,$$
(3)

where where for all $\delta(i, a_i, a'_i, a_{-i}) = (u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i}))$, for all a_i , for all a'_i , and α_{-i} shows that π is restricted by a finite collection of linear constraints in a finitely dimensional equilidean space. Thus the set of correlated equilibria is convex.

• This implies that also the utility set induced by correlated equilibria is a convex. This can be handy in applications. The problem of maximizing

joint surplus in a social activity by some contractual arrangement reduces to a linear programming problem: maxmimize surplus under the constraint