Microeconomics

2. Game Theory

Alex Gershkov

http://www.econ2.uni-bonn.de/gershkov/gershkov.htm

21. November 2007

Strategic form games

1.a Describing a game in strategic form

1.b Iterated deletion of strictly dominated strategies

1.c Nash equilibrium and examples

1.d Mixed strategies

1.e Existence of Nash equilibria

1.f Rationilazability

1.g Correlated equilibrium

1.a Describing a game in strategic form

Example: Entry game

1. 2 ice cream vendors decide whether or not to open an ice cream stand in a particular street
2. the decision is taken without observing the other vendor's action
3. if only one stand is opened, this vendor earns \$ 3 (the other vendor earns zero)
4. if both stands open, they share profits and earn \$ 1.5 each

1.a Describing a game in strategic form

Let's use a matrix to organise this information

	entry	*no entry*
entry	1.5,1.5	3,0
no entry	0,3	0,0

we use the following conventions

1. rows contain vendor 1's decisions, columns contain vendor 2's decisions
2. each outcome of the game is located in one cell of the matrix
3. outcomes (payoffs, profits) are given in expected utility in the form (u_1, u_2).

1.a Describing a game in strategic form

This way of describing 'simultaneous' move games is called strategic form (sfg). Formally, a sfg consists of

1. a (not necessarily finite) set of players N: in the example $N = \{1, 2\}$, generally $N = \{1, \ldots, n\}$
2. a (not necessarily finite) set of *pure* strategies $S = S_1 \times S_2$: in the example $S_i = \{entry, no\ entry\}$, $S = S_i^2,\ i \in \{1, 2\}$, generally $S = \{S_1 \times \ldots \times S_n\} = S_i^n$
3. a set of payoff fns $u(S)$: given as discrete values by the payoffs in the sample matrix, generally by a vector of *expected utility* fns $u(S) = u_1(S), \ldots, u_n(S)$

Thus a game in strategic form is fully described by $\{N, S, u\}$. A sfg is used to describe a game with no time dimension.

1.a Describing a game in strategic form

Predicting the agent's (=vendor's) actions is easy provided that

1. payoffs (=profits) are the only thing agents care about
2. agents are rational
3. agents maximise their profits

Then the predicted actions are (*entry*,*entry*) leading to an outcome of (1.5,1.5).

Notice that in this example an agent does not need to know the other agent's choice in order to determine his optimal choice.

1.b Equilibria (idsds)

Example:

	l	c	r
t	0,0	4,-1	1,-1
m	-1,4	5,3	3,2
b	-1,2	0,1	4,1

1.b Equilibria (idsds)

Definition: Given player i's pure strategy set S_i, a *mixed strategy* for player i, $\sigma_i : S_i \longrightarrow [0, 1]$ is a probability distribution over pure strategies. (Denote by Σ_i the space of player i mixed strategies and $\Sigma = \Sigma_1 \times \ldots \times \Sigma_n$)

Another example:

	l	r
t	2,0	-1,1
m	0,1	0,0
b	-1,0	2,2

1.b Equilibria (idsds)

Definition: Pure strategy s_i is *strictly dominated* for player i if there exists $\sigma'_i \in \Sigma_i$ such that

$$u_i(\sigma'_i, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i} \tag{1}$$

where $S_{-i} = S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$

The strategy s_i is *weakly dominated* if there exists a $\sigma'_i \in \Sigma_i$ such that inequality (1) holds weakly, that is, if the inequality is strict for at least one s_{-i} and equality holds for at least one s_{-i}.

Definition: A <u>rational</u> player should never play a strictly dominated strategy.

1.b Equilibria (idsds) & Rationality

Definition: X is called common knowledge (ck) between players A and B if for any $i \in \{A, B\}$, $i \neq j$, $k = 1, 2, 3, \ldots$ $\underbrace{K_i K_j K_i \ldots}_{k\times} X$

Idsds requires the following assumptions

1. players are rational
2. common knowledge of rationality
3. players know their payoffs
4. common knowledge of the structure of the game.

1.b Equilibria (idsds)

Definition

The process of idsds proceeds as follows: Set $S_i^0 = S_i$ and $\Sigma_i^0 = \Sigma_i$. Define S_i^k recursively by

$$S_i^k = \{s_i \in S_i^{k-1} | \text{ there is no } \sigma_i \in \Sigma_i^{k-1}$$

$$\text{s.t. } u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}^{k-1}\}$$

and define

$$\Sigma_i^k = \{\sigma_i \in \Sigma_i | \sigma_i(s_i) > 0 \text{ only if } S_i \in S_i^k\}$$

The set S_i^∞ (Σ_i^∞) is the set of Pi's pure (mixed) strategies that survive idsd.

1.c Nash equilibrium

Game ('Battle of the sexes')

	f	o
f	1,2	0,0
o	0,0	2,1

Here idsds does not work—we need a sharper tool: NE'qm.

NE'qm puts an additional restriction on the knowledge of players compared to idsds: It assumes that players have a correct expectation of which prediction is played and subsequently have no incentive to deviate from this prediction!

1.c Nash equilibrium

Definition: A strategy profile s is a vector of dimension n containing a single strategy choice for each player.

Notation: $s = (s_1, \ldots, s_n) = (s_i, s_{-i})$

Definition: Player i's best response $\mathcal{B}_i$ is the set of own strategies giving the highest payoff when the opponents play s_{-i}. Formally

$$\mathcal{B}_i(s_{-i}) = \{s_i \in S_i | u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}), \ \forall s_i' \in S_i\}$$

1.c Nash equilibrium

Definition: A Nash equilibrium is a strategy profile s^* which prescribes a best response s_i^* for each player i given that the opponents play s_{-i}^*. Hence for every $i \in N$

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*),\ \forall s_i' \in S_i$$

or equivalently

$$s_i^* \in \mathcal{B}_i(s_{-i}^*),\ \forall i \in N.$$

Notice that according to this definition, the player's expected prediction of play is correct in equilibrium.

1.c Nash equilibrium

Consider the following game

	l	r
u	3,2	2,0
d	0,0	1,1

Observe that

$$\mathcal{B}_1(l) = \{u\},\ \mathcal{B}_1(r) = \{u\}$$

while

$$\mathcal{B}_2(u) = \{l\},\ \mathcal{B}_2(d) = \{r\}.$$

Hence the unique NE'qm is (u,l).

1.d NE'qm & mixed strategies

What about this game, known as *matching pennies*?

	h	t
h	1,-1	-1,1
t	-1,1	1,-1

THIS GAME HAS NO PURE STRATEGY NE'QM!

Rem: A mixed strategy σ_i is a probability distribution over player i's set of pure strategies S_i denoted by Σ_i.

Assumption

We assume that i randomises independently from the other players' randomisations.

Remark

Σ_i *is an* $|S_i|$*-dimensional simplex* $\Sigma_i = \{\mathbf{p} \in [0,1]^{|S_i|} : \sum p_h = 1\}$.

1.d NE'qm & mixed strategies

Definition: The mixed extension of the sfg $\{N, S, u\}$ is the sfg $\{N, \Sigma, u\}$ where $i's$ payoff fn $u_i : \times_{j \in N} \Sigma_j \to \mathbb{R}$ assigns to each $\sigma \in \times_{j \in N} \Sigma_j$ the expected utility of the lottery over S induced by σ.

Definition: A mixed strategy NE'qm of a sfg is a NE'qm of its mixed extension such that for every $i \in N$

$$\sigma_i^* \in \operatorname*{argmax}_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \sigma_{-i}).$$

Lemma

If σ^* is a mixed strategy NE'qm then for every $i \in N$ every pure action $s_i \in S_i$ to which σ_i^* attaches positive probability is a best response to σ_{-i}^*.

1.d NE'qm & mixed strategies

It is sufficient to check that no player has a profitable pure-strategy deviation.

Alternative definition: A mixed strategy profile σ^* is a NE'qm if for all players i

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

1.d NE'qm & mixed strategies

Theorem

(Nash 1950) Every finite sfg $\{N, S, u\}$ *has a mixed strategy Nash equilibrium.*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*), \; \forall \sigma_i \in \Sigma_i.$$

Proof: The proof consists of showing that our setting satisfies all of Kakutani's assumptions and then concluding that there exists a σ^* **which is contained** in the set of best responses $\mathcal{B}(\sigma^*)$.

1.e Existence of NE'qm

The existence question boils down to asking when the set of intersections of best response correspondences

$$\mathcal{B}_i(\sigma_{-i}) = \{\sigma_i \in \Sigma_i | u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}),\ \forall \sigma_i' \in \Sigma_i\}$$

is non-empty. The way Nash answered this question was to set up a mapping $\mathcal{B} : \Sigma \to \Sigma$ where

$$\mathcal{B}(\sigma) = (\underbrace{\mathcal{B}_1(\sigma_{-1})}_{\sigma_1}, \ldots, \underbrace{\mathcal{B}_n(\sigma_{-n})}_{\sigma_n}).$$

1.e Existence of NE'qm

Definition

Given $X \subset \mathbb{R}^n$, a **correspondence** F between the sets X and $Y \subset \mathbb{R}^k$, written $F : X \rightrightarrows Y$, is a mapping that assigns a set $F(x)$ to every $x \in X$.

Definition

The **graph** of a correspondence $F : X \rightrightarrows Y$ is the set $\Gamma(F) = \{(x, y) \in X \times Y \mid y \in F(x)\}$.

Hence the graph $\Gamma(F)$ is a set which can be checked for open or closedness just like any other set.

1.e Existence of NE'qm

Theorem

(Kakutani 1941)

A correspondence $F : X \rightrightarrows X$ *satisfying*

1. X *is a convex, closed, and bounded subset of* $\mathbb{R}^n$,
2. $F(\cdot)$ *is non-empty and convex-valued, and*
3. $F(\cdot)$ *has a closed graph*

has a fixed point $x^* \in F(x^*)$ *for* $x^* \in X$.

1.e Existence of NE'qm

1. Σ_i is a convex, closed, and bounded subset of $\mathbb{R}^n$:
   0.1 Σ_i is convex because if $\sigma_i,\ \sigma_i' \in \Sigma_i$ then any mixture $\lambda\sigma_i + (1-\lambda)\sigma_i' \in \Sigma_i$ for all $\lambda \in [0,1]$.
   0.2 Σ_i is closed because its boundary points (assigning probability 0 and 1 to some pure strategy) are included in the set.
   0.3 Σ_i is bounded because all probabilities are bounded between $[0,1]$.

1.e Existence of NE'qm

2. $\mathcal{B}_i(\sigma_{-i})$ is non-empty and convex-valued:
   0.1 There is a finite number of pure strategies (by assumption), therefore there exists at least one best reply to a given profile. Thus the best response correspondence cannot be empty.
   0.2 Let $\sigma_i, \ \sigma_i' \in \Sigma_i$ be best replies, then the mixture is a best reply as well because it will yield the same amount. Hence, $\mathcal{B}_i(\sigma_{-i})$ is convex valued.

1.e Existence of NE'qm

3. $\mathcal{B}(\cdot)$ has a closed graph; thus any boundary point of the graph is included.

   Assume the condition is violated, so there is a sequence $(\sigma^k, \widehat{\sigma}^k) \longrightarrow (\sigma, \widehat{\sigma})$ with $\widehat{\sigma}^k \in \mathcal{B}(\sigma^k)$, but $\widehat{\sigma} \notin \mathcal{B}(\sigma)$. Then, $\widehat{\sigma}_i \notin \mathcal{B}_i(\sigma_{-i})$ for some player i. Thus, there is an $\varepsilon > 0$ and a σ_i' such that $u_i(\sigma_i', \sigma_{-i}) > u_i(\widehat{\sigma}_i, \sigma_{-i}) + 3\varepsilon$. Since u_i is continuous and $(\sigma^k, \widehat{\sigma}^k) \longrightarrow (\sigma, \widehat{\sigma})$, for k sufficiently large we have

   $$u_i(\sigma_i', \sigma_{-i}^k) > u_i(\sigma_i', \sigma_{-i}) - \varepsilon > u_i(\widehat{\sigma}_i, \sigma_{-i}) + 2\varepsilon > u_i(\widehat{\sigma}_i^k, \sigma_{-i}^k) + \varepsilon.$$

   Therefore, σ_i' does strictly better against σ_{-i}^k than $\widehat{\sigma}_i^k$ does, contradiction to $\widehat{\sigma}^k \in \mathcal{B}(\sigma^k)$.

1.f Rationalisability

Bernheim's (1984) and Pearce's (1984) concept of Rationalisability stems from the desire to base an equilibrium concept on purely behavioural assumptions (rationality)

- agents view their opponents' choice as an uncertain events
- every agent optimizes s.t. some probabilistic assessment of uncertain event
- this assessment is consistent with all of his information
- the previous points are commonly known

The starting point is: What are *all* the strategies that a rational player could play? A rational player will use only those strategies that are a best response to some belief he might have about the strategies of his opponents.

Definition

A player's belief $\mu_i(\sigma_j)$ is the conditional probability he attaches to opponent j playing a particular strategy σ_j.

1.f Rationalisability

Consider the following game due to Bernheim (1984)

	b_1	b_2	b_3	b_4
a_1	0,7	2,5	7,0	0,1
a_2	5,2	3,3	5,2	0,1
a_3	7,0	2,5	0,7	0,1
a_4	0,0	0,-2	0,0	10,-1

1. (a_2, b_2) is the only Nash equilibrium (and hence both a_2 and b_2 are rationalizable)
2. strategies a_1, a_3, b_1 and b_3 are rationalizable as well
3. a_4 and b_4 are **not** rationalizable

1.f Rationalisability

Definition

Set $\widetilde{\Sigma}_i^0 \equiv \Sigma_i$ and for each i recursively

$$\widetilde{\Sigma}_i^k \equiv \{\sigma_i \in \widetilde{\Sigma}_i^{k-1} | \exists \sigma_{-i} \in \times_{j \neq i} \widetilde{\Sigma}_j^{k-1}$$

$$\text{s.t. } u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \text{ for all } \sigma_i' \in \widetilde{\Sigma}_i^{k-1}\}$$

The strategy σ_i of i is *rationalizable* if $\sigma_i \in \bigcap_{k=0}^{\infty} \widetilde{\Sigma}_i^k$.

A strategy profile σ is *rationalizable* if σ_i is rationalizable for each player i.

1.f Rationalisability and idsd

Clearly, with any number of players, a strictly dominated strategy is never a best response: If σ'_i strictly dominates σ_i relative to Σ_{-i}, then σ'_i is a strictly better response than σ_i to every σ_{-i} in Σ_{-i}.

Thus, the set of rationalizable strategies is contained in the set that survives iterated strict dominance.

Theorem

(Pearce 1984) The set of rationalizable strategies and the set that survives idsds coincide in two-players games.

1.f Rationalisability

Proof: (Second direction) X - the set of strategies that survive idsd.

We show that for every player every member of X_j is rationalizable. By definition, no action in X_j is strictly dominated in the game with the set of actions for each player i is X_i. $\Rightarrow$ every action in X_j is a BR among X_j to some beliefs on X_{-j}.

We have to show that every action in X_j is a BR among S_j to some beliefs on X_{-j}. If $s_j \in X_j$ is not a BR among S_j then there is k s.t. s_j is a BR among X_j^k to a belief μ_j on X_{-j}, but is not a BR among X_j^{k-1}. Then, $\exists b_j \in X_j^{k-1} \backslash X_j^k$ that is a BR among X_j^{k-1} to μ_j, contradiction, since b_j was eliminated at n.

1.h Correlated equilibrium

NE'qm assumes that players use independent randomisation in order to arrive at their mutual actions. Given these, NE'qm is the minimal necessary condition for reasonable predictions (in general games). What if we dispense with this independence?

Consider that players can engage in (unmodelled) pre-play communication. Then they could conceivably decide to choose their actions based on this communication or, more generally, on any other public (or private) event.

In order to capture this idea, let's assume that the players can construct (or have access to) a signalling device which sends (public or private) signals on which the players can co-ordinate their actions.

1.h Correlated equilibrium

Consider the following example

	l	r
u	5,1	0,0
d	4,4	1,5

This game has 2 pure NE'qa $\{(u,l),(d,r)\}$ plus the mixed e'qm $((^1/_2)u,(^1/_2)l)$ leading to $(2.5,2.5)$.

Can the players do any better?

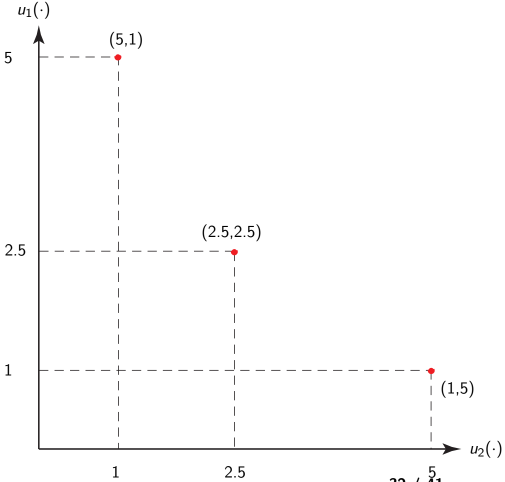

1.h Correlated equilibrium

Suppose players can agree on constructing a public randomisation device with 2 equally likely outcomes $\{\mathcal{H}, \mathcal{T}\}$.

Assume that the publicly observable outcomes of the randomisations of this device are interpreted as the following plan c:

- if $\mathcal{H}$ (happening with $\text{pr}(\mathcal{H}) = {}^1\!/_2$), then play (u, l)
- if $\mathcal{T}$ (happening with $\text{pr}(\mathcal{T}) = {}^1\!/_2$), then play (d, r)

Then the expected payoffs from this plan c are

$$u_1(c) = \frac{1}{2}u_1(u, l) + \frac{1}{2}u_1(d, r) = \frac{1}{2}5 + \frac{1}{2}1 = 3 > 2.5!$$

(P2's $u_2(\cdot)$ is symmetric.)

1.g Correlated equilibrium

In the same fashion, by using their public randomisation device to co-ordinate their actions, the players can do better than Nash by attaining any point in the convex hull of NE'qa!

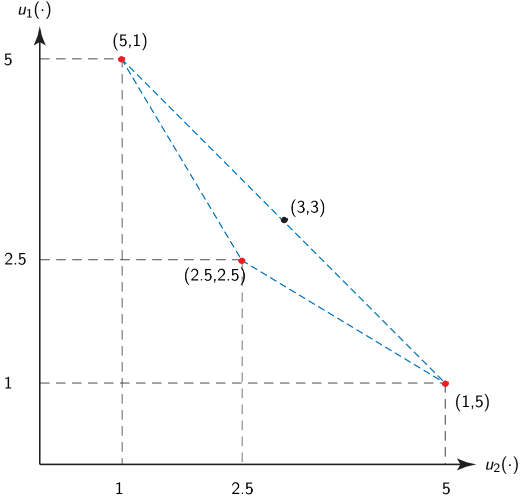

1.h Correlated equilibrium

Can we do still better?

Suppose that players can agree on establishing a device which sends correlated but private signals to each player.

Consider a device with 3 equally likely outcomes $\{\mathcal{A}, \mathcal{B}, \mathcal{C}\}$ and assume the following information partitions:

- P1's partition $\mathcal{P}_1 = \{\{\mathcal{A}\}, \{\mathcal{B}, \mathcal{C}\}\}$
- P2's partition $\mathcal{P}_2 = \{\{\mathcal{A}, \mathcal{B}\}, \{\mathcal{C}\}\}$

Let's look at the following plan d:

- P1 plays u if told $\{\mathcal{A}\}$ and plays d if told $\{\mathcal{B}, \mathcal{C}\}$
- P2 plays l if told $\{\mathcal{A}, \mathcal{B}\}$ and plays r if told $\{\mathcal{C}\}$

1.h Correlated equilibrium

Assumption

Players are Bayesian rational, ie. they use Bayes' rule to update their beliefs.

Claim: No player has an incentive to deviate from d.

Now P1's overall payoff from using the plan d is

$$u_1(d) = \frac{1}{3}u_1(u,l) + \frac{1}{3}u_1(d,l) + \frac{1}{3}u_1(d,r) = \frac{1}{3}5 + \frac{1}{3}4 + \frac{1}{3}1 = \frac{10}{3} > 3!$$

(Again the case for P2 is symmetric.)

1.h Correlated equilibrium

By using a <u>private</u> randomisation device, even payoffs <u>outside</u> the convex hull of NE'qa are attainable!

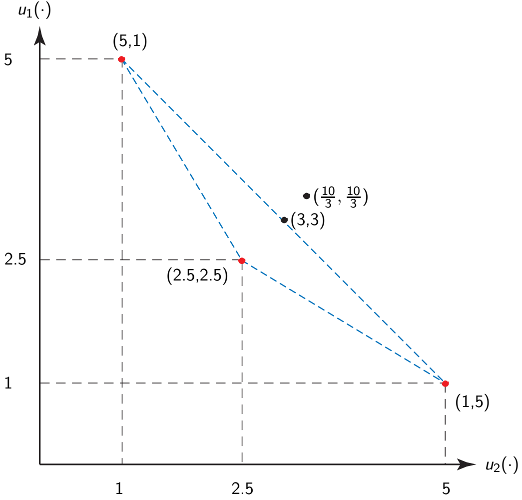

1.h Correlated equilibrium

Definition

Communication device is a triple (Ω, H_i, p) where

- Ω is a state space (set of possible outcomes of the device)
- p is a probability measure on Ω
- H_i is information partition for i, if the true state is ω, Pi is told that the state lies in $h_i(\omega)$, Pi's posterior about ω is given by Bayes' rule $p(\omega|h_i) = p(\omega)/p(h_i)$ for $\omega \in h_i$ and 0 for $\omega \notin h_i$

Definition

A pure strategy for the expended game is $\rho_i : H_i \longrightarrow S_i$

Alternatively,

A pure strategy is $\rho_i : \Omega \rightarrow S_i$, such that $\rho_i(\omega) = \rho_i(\omega')$ whenever $\omega, \omega' \in h_i$ for some $h_i \in H_i$

1.h Correlated equilibrium

Definition

A CE'qm ς relative to information structure (Ω, H_i, p) is a Nash equilibrium in strategies that are adapted to this information structure. That is. $(\varsigma_1, ..., \varsigma_N)$ is a correlated equilibrium if, for every i and every adapted strategy $\widehat{\varsigma}_i$

$$\sum_{\omega \in \Omega} p(\omega) u_i(\varsigma_i(\omega), \varsigma_{-i}(\omega)) \geq \sum_{\omega \in \Omega} p(\omega) u_i(\widehat{\varsigma}_i(\omega), \varsigma_{-i}(\omega)).$$

1.h Correlated equilibrium

Theorem

(Aumann 1974) For every finite game, the set of CE'qa is non-empty.

Proof: Let the recommendation be an independent randomisation for every player. Than the CE'qm is an independently mixed NE'qm which we know to exist. Hence the set of CE'qa contains the set of NE'qa.

Theorem

(Aumann 1974) The set of CE'qm payoffs of a sfg is convex.

Proof: One can always obtain a new CE'qm from publicly randomising between two other CE'qa.

1.h Correlated equilibrium

Since players have access to a public randomisation device, the set of CE'qa is convex (and contains the set of NE'qa)!