

Advanced microeconomics 3: Game Theory

Slide set 1: Basic concepts of game theory

Pauli Murto

Aalto University

January 26, 2023

Outline

- This slide set covers the two first weeks of the course
- We go through the basic elements and solution concepts of non-cooperative game theory
- Reading material for these lectures:
 - Mas-colell, Whinston, Green: Ch. 7, Ch. 8 A - D.
 - Mailath: Ch. 1, 2.1, 2.4 - 2.6, 4.1
- Other relevant sources:
 - Osborne and Rubinstein: Ch. 1 - 4, 6.1
 - Fudenberg and Tirole: Ch. 1 - 3.4
 - Myerson: Ch. 1, 2.1-2.5, 3.1-3.8
 - Maschler, Solan, Zamir: Ch. 1-6

Background

- In Microeconomics 1 and 2, the focus was on individual decision making
- Collective decisions appeared in the context of general equilibrium:
 - Every agent optimizes against the prices
 - Prices equate supply and demand
 - No need for the agents to understand where prices come from
- Microeconomics 3 and 4 will be explicitly concerned about strategic interaction between the agents
- This requires more sophisticated reasoning from agents: how are the other players expected to behave?
- Microeconomics 3 studies the methodology of modeling interactive decision making: *game theory*

Some classifications

- Non-cooperative vs. cooperative game theory
 - In non-cooperative games, individual players and their optimal actions are the primitives
 - In cooperative games, coalitions of players and their joint actions are the primitives
 - In this course, we concentrate on non-cooperative games
- Static vs. dynamic games
 - We start with static games but move then to dynamic games
- Games with complete vs. incomplete information
 - We start with games with complete information but then move to incomplete information

From decision theory to game theory

- We maintain the decision theoretic framework familiar from Part I of the course
- In particular, the decision maker is *rational*:
 - Let A denote the set of available actions
 - An exhaustive set of possible consequences C .
 - A consequence function $g : A \rightarrow C$ specifying which actions lead to which consequences.
 - Complete and transitive preference relation $\succsim$ on C .
- Preference relation can be represented by a real valued *utility function* u on C

- To model decision making under *uncertainty*, we assume that the preferences also satisfy von Neumann -Morgenstern axioms.
- This leads to expected utility maximization:
 - Let the consequence depend not only decision maker's action but also a state in some set Ω .
 - If the decision maker takes action a , and state $\omega \in \Omega$ materializes, then the decision maker's payoff is $u(g(a, \omega))$.
 - If the uncertainty is captured by a probability distribution p on Ω , the decision maker maximizes his expected payoff

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

- In decision theory uncertainty in the parameters of environment or events that take place during the decision making process
- In strategic situations, consequence g not only depends on the DM's own action, but also on the actions of the other DMs
- Hence, uncertainty may stem from random actions of other players or from reasoning of other players

A strategic form game

- A game in strategic (or normal) form consists of:
 - 1 Set $\mathcal{I} = \{1, \dots, I\}$ of players
 - 2 Pure strategy space S_i for each $i \in \mathcal{I}$
 - 3 A von Neumann-Morgenstern utility u_i for each $i \in \mathcal{I}$:

$$u_i : S \rightarrow \mathbb{R},$$

where $S := \times_{i=1}^I S_i$.

- That is, $u_i(s)$ gives the utility of i for strategy profile $s := (s_1, \dots, s_I)$.

- We write $u := (u_1, \dots, u_I)$
- We also write $s = (s_i, s_{-i})$, where $s_{-i} \in S_{-i} := \times_{j \neq i} S_j$
- We often (but not always) assume that S_i are finite sets.
- The game is hence defined by $\langle \mathcal{I}, \{S_i\}_{i \in \mathcal{I}}, \{u_i\}_{i \in \mathcal{I}} \rangle$
- In standard table presentation, player 1 chooses row and player 2 chooses column. Each cell corresponds to payoffs so that player 1 payoff is given first.
- Some classic 2x2 games that highlight different strategic aspects:

Prisoner's dilemma:

	<i>Cooperate</i>	<i>Defect</i>
<i>Cooperate</i>	3, 3	0, 4
<i>Defect</i>	4, 0	1, 1

Coordination game ("stag hunt"):

	A	B
A	2, 2	0, 1
B	1, 0	1, 1

Battle of sexes:

	<i>Ballet</i>	<i>Football</i>
<i>Ballet</i>	2, 1	0, 0
<i>Football</i>	0, 0	1, 2

Hawk-Dove:

	<i>Dove</i>	<i>Hawk</i>
<i>Dove</i>	3, 3	1, 4
<i>Hawk</i>	4, 1	0, 0

Matching pennies:

	<i>Head</i>	<i>Tail</i>
<i>Head</i>	1, -1	-1, 1
<i>Tail</i>	-1, 1	1, -1

Mixed strategies

- The players may also randomize their actions, i.e. use mixed strategies
- Suppose that strategy set S_i is finite: $S_i = \{s_{i1}, \dots, s_{in}\}$

Definition

A mixed strategy for player i , $\sigma_i : S_i \rightarrow [0, 1]$ assigns to each pure strategy $s_{ik} \in S_i$, $k = 1, \dots, n$ a probability $\sigma_i(s_{ik}) \geq 0$ that it will be played, such that

$$\sum_{k=1}^n \sigma_i(s_{ik}) = 1.$$

- The mixed strategy space for i is a simplex over the pure strategies

$$\Delta(S_i) := \left\{ (\sigma_i(s_{i1}), \dots, \sigma_i(s_{in})) \in \mathbb{R}^n : \sigma_i(s_{ik}) > 0 \ \forall k, \sum_{k=1}^n \sigma_i(s_{ik}) = 1 \right\}.$$

- $\sigma := (\sigma_1, \dots, \sigma_I)$ is a mixed strategy profile
- If the players choose their strategies simultaneously and independently of each other, a given pure strategy profile $(s_1, \dots, s_I)$ is chosen with probability

$$\prod_{i=1}^I \sigma_i(s_i).$$

- Player i 's payoff to profile σ is

$$u_i(\sigma) = \sum_{s \in S} \left(\prod_{i=1}^I \sigma_i(s_i) \right) u_i(s).$$

- Note that here we utilize the von Neumann - Morgenstern utility representation
- Mixed strategies over continuous pure strategy spaces are defined analogously
- The game $\langle \mathcal{I}, \{\Delta(S_i)\}_{i \in \mathcal{I}}, \{u_i\}_{i \in \mathcal{I}} \rangle$ is sometimes called the mixed extension of the game $\langle \mathcal{I}, \{S_i\}_{i \in \mathcal{I}}, \{u_i\}_{i \in \mathcal{I}} \rangle$

Extensive form

- Strategic form seems to miss some essential features of strategic situations: dynamics and information
- We must specify:
 - Who moves when?
 - What do players know when they move?
 - What are payoffs under all contingencies?

Example: a simple card game

- Players 1 and 2 put one dollar each in a pot
- Player 1 draws a card from a stack and observes it privately
- Player 1 decides whether to raise or fold
- If fold, then game ends, and player 1 takes the money if the card is red, while player 2 takes the money if black
- If raise, then player 1 adds another dollar in the pot, and player 2 must decide whether to meet or pass
- If player 2 passes, the game ends and player 1 takes the money in the pot
- If player 2 meets, he adds another dollar in the pot. Then player 1 shows the card, and the game ends. Again, player 1 takes the money in the pot if the card is red, while player 2 takes the money if black

A finite game in extensive form consists of:

1. The set of players, $\mathcal{I} = \{1, \dots, I\}$.
2. A directed graph i.e. a set nodes X and arrows connecting the nodes. This must form a tree which means:
 - There is a single initial node x^0 , i.e. a node with no arrows pointing towards it.
 - For each node, there is a uniquely determined path of arrows connecting it to the initial node. (This is called the path to the node).
3. The nodes are divided into:
 - Terminal nodes Z , i.e. with no outward pointing arrows.
 - Decision nodes $X \setminus Z$, i.e. nodes with outward pointing arrows.

- Each decision node is labeled as belonging to a player in the game (the player to take the decision). This labeling is given by a function $\iota : X \setminus Z \rightarrow \mathcal{I}$.
- Each arrow represents an action available to the player at the decision node at the origin of the arrow. Actions available at node x is $A(x)$. If there is a path of arrows from x to x' , then we say that x' is a successor of x , and we write $x' \in s(x)$.
- Payoffs assign a utility number to each terminal payoff (and thus also to each path through the game tree). Each player i has a payoff function $u_i : Z \rightarrow \mathbb{R}$.

7. A partition H of decision nodes (i.e. $H = (h^1, \dots, h^K)$ such that $h^k \subset X \setminus Z$ for all k and $h^k \cap h^l = \emptyset$ for $k \neq l$ and $\cup_k h^k = X \setminus Z$) into information sets h^k . These are collections of nodes such that:
- ① The same player acts at each node within the information set. (i.e. $\iota(x) = \iota(x')$ if $\exists k$ such that $x, x' \in h^k$).
 - ② The same actions must be available at all nodes within the information set. (i.e. $A(x) = A(x')$ if $\exists k$ such that $x, x' \in h^k$).
8. If Nature moves, the probability that she takes each available action must be specified.

Remarks:

- Simultaneous actions can be modeled by an extensive form, where one player moves first, but so that all nodes resulting from her actions are in a single information set for the other player
- Asymmetric information can be modeled by moves by nature and appropriately chosen information sets (in particular, Bayesian games, to be analyzed later)

Some classifications

- Games with perfect information
 - If all information sets are singletons, then a game has perfect information.
 - Otherwise the game has imperfect information.
 - Note that since only one player moves in each node, games of perfect information do not allow simultaneous actions
- Multi-stage games with observed actions
 - There are "stages" $k = 1, 2, \dots$ such that
 - 1 In each stage k every player knows all the actions taken in previous stages (including actions taken by Nature)
 - 2 Each player moves at most once within a given stage
 - 3 No information set contained in stage k provides information about play in that stage
 - In these games, all actions taken before stage k can be summarized in public history h^k

Some classifications

- Bayesian games, or games of incomplete information
 - Nature chooses a "type" for each player according to a common prior
 - Each player observes her own type but not that of others
 - These games will be extensively analyzed towards the end of the course
- Games of perfect recall
 - A game is of perfect recall, when no player forgets information that he once knew
 - A formal definition involves some restrictions on information sets
 - All the games that we will consider are of perfect recall

Strategies

- Let H_i be the set of player i 's information sets, and let $A(h_i)$ be the set of actions available at $h_i \in H_i$
- The set of all actions for i is then $A_i := \cup_{h_i \in H_i} A(h_i)$

Definition

A pure strategy for i is a map

$$s_i : H_i \rightarrow A_i$$

with $s_i(h_i) \in A(h_i)$ for all $h_i \in H_i$.

Strategies

- Important: a strategy must define the action for i at *all* contingencies defined in the game
- The set of pure strategies for i is

$$S_i = \prod_{h_i \in H_i} A(h_i).$$

In a finite game, this is a finite set.

Mixed strategies

- Having defined pure strategies, we can define mixed strategies just as in the case of strategic form:

Definition

A mixed strategy for player i , $\sigma_i : S_i \rightarrow [0, 1]$ assigns to each pure strategy $s_i \in S_i$ a probability $\sigma_i(s_i) \geq 0$ that it will be played, such that

$$\sum_{s_i \in S_i} \sigma_i(s_i) = 1.$$

Behavior strategies

- There is another more convenient way to define mixed strategies, called behavior strategies.
- With those strategies, mixing takes place independently at each decision node:

Definition

A behavior strategy for player i specifies for each $h_i \in H_i$ a probability distribution on the set of available actions $A(h_i)$. That is,

$$b_i \in \prod_{h_i \in H_i} \Delta(A(h_i)),$$

where $\Delta(A(h_i))$ is a simplex over $A(h_i)$.

- Every mixed strategy generates a unique behavior strategy (see e.g. Fudenberg-Tirole section 3.4.3 for the construction)
- In games of perfect recall (all relevant games for our purpose), it makes no difference whether we use mixed or behavior strategies:

Theorem (Kuhn 1953)

In a game of perfect recall, mixed and behavior strategies are essentially equivalent.

- More precisely: every mixed strategy is equivalent to the unique behavior strategy it generates, and each behavior strategy is equivalent to every mixed strategy that generates it.
- Therefore, there is no loss in using behavior strategies

From extensive form to strategic form

- Recall that extensive form defines payoff $u_i : Z \rightarrow \mathbb{R}$ for each terminal node.
- Since each strategy profile leads to a probability distribution over terminal nodes Z , we may directly associate payoffs for strategy profiles (utilizing expected utility formulation):

$$u_i : S \rightarrow \mathbb{R},$$

where $S := \times_{i=1}^I S_i$.

- Now $\langle \mathcal{I}, \{S_i\}_{i \in \mathcal{I}}, \{u_i\}_{i \in \mathcal{I}} \rangle$ meets our definition of a strategic form game
- This is the strategic-form representation of our extensive form game

- To see how this works, take any 2×2 game, formulate its extensive form assuming sequential moves, and then move back to strategic form (and you get a 2×4 game)
- Every extensive form game may be represented in strategic form
- However, as will be made clear later, we will need extensive form to refine solution concepts suitable for dynamic situations

Solution concepts

- We now develop the basic solution concepts for the strategic form games
- This is the simple game form normally used for analysing static interactions
- But any extensive form game may be represented in strategic form, so the concepts that we develop here apply to those as well

Implications of rationality

- Rationality means that each of the players chooses s_i in order to maximize her expectation of u_i
- But what should players expect of other player's actions?
- Standard game theory typically assumes that the rationality of the players is *common knowledge*.
- This means that all the players are rational, all the players know that all the players are rational, all the players know that all the players know that all the players are rational, and so on...

Implications of rationality

- We start by asking what implication *rationality* has on the players' behavior
- We then ask what *common knowledge of rationality* implies for the players' behavior
- This will lead us to the concept of rationalizable strategies

Dominant strategies

Let us start with the most straight-forward concepts:

Definition

A strategy s_i is a dominant strategy for i if for all $s_{-i} \in S_{-i}$ and for all $s'_i \neq s_i$,

$$u_i(s_i, s_{-i}) > u(s'_i, s_{-i}).$$

Definition

A strategy s_i is a weakly dominant strategy for i if for all $s'_i \neq s_i$,

$$u_i(s_i, s_{-i}) \geq u(s'_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i} \text{ and}$$

$$u_i(s_i, s_{-i}) > u(s'_i, s_{-i}) \text{ for some } s_{-i} \in S_{-i}.$$

- In a few cases, this is all we need.

Example: Prisoner's dilemma

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

	C	D
C	3, 3	0, 4
D	4, 0	1, 1

- Whatever strategy the other player chooses, it is strictly optimal for i to choose D and not C . Thus (D, D) is the *dominant strategy equilibrium* of this game.
- Thus, rational players should always play D (even if (C, C) would be better for both)

Example: Second-price auction

- A seller has one indivisible object for sale
- There are I potential buyers with valuations $0 \leq v_1 \leq \dots \leq v_I$, and these valuations are common knowledge
- The bidders simultaneously submit bids $s_i \geq 0$.
- The highest bidder wins the object and pays the second highest bid (if several bidders bid the highest price, then the good is allocated randomly among them)
- Exercise: show that for each player i bidding $s_i = v_i$ weakly dominates all other strategies
- Thus, $s_i = v_i$ for all i is a *weakly dominant strategy equilibrium*
- Bidder I wins and has payoff $v_I - v_{I-1}$.

Dominated strategies

- However, very few games have dominant strategies for all players
- Consider the following game:

	<i>L</i>	<i>M</i>	<i>R</i>
<i>U</i>	4, 3	5, 1	6, 2
<i>M</i>	2, 1	8, 4	3, 6
<i>D</i>	3, 0	9, 6	2, 8

- There are no dominant strategies, but M is *dominated* by R , thus a rational player 2 should not play M
- But if player 2 will not to play M , then player 1 should play U
- But if player 1 will play U , player 2 should play L
- This process of elimination is called iterated strict dominance
- We say that a game is solvable by iterated strict dominance when the elimination process leaves each player with only a single strategy

- Note that a pure strategy may be strictly dominated by a mixed strategy even if not dominated by a pure strategy. Below, M is not dominated by U or D , but it is dominated by playing U with prob. $1/2$ and D with prob. $1/2$:

	L	R
U	2, 0	-1, 0
M	0, 0	0, 0
D	-1, 0	2, 0

Definition

Pure strategy s_i is strictly dominated for player i if there exists $\sigma_i \in \Delta(S_i)$ such that

$$u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}$$

Iterated deletion of dominated strategies

Definition

The process of iterated deletion of strictly dominated strategies proceeds as follows: Set $S_i^0 = S_i$. Define S_i^n recursively by

$$S_i^n = \{s_i \in S_i^{n-1} \mid \nexists \sigma_i \in \Delta(S_i^{n-1}) \text{ s.t. } u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}^{n-1}\}$$

Set

$$S_i^\infty = \bigcap_{n=0}^\infty S_i^n.$$

S_i^∞ is the set of player i 's pure strategies that survive iterated deletion of strictly dominated strategies.

- The game is solvable by iterated (strict) dominance if, for each player i , S_i^∞ is a singleton.
- Strict dominance is attractive since it is directly *implied by common knowledge of rationality*: rational players never use strictly dominated strategies, hence common knowledge of rationality suggests that players should not use strategies that are eliminated in the iterated process described above.
- In process defined here, one deletes simultaneously all dominated strategies for both players in each round. One can show that the details of elimination process do not matter.
- One can also apply iterative deletion to weakly dominated strategies, but then the order of deletion matters

Example: Cournot model with linear demand (1)

- Let us model the two-firm Cournot model as a game $\langle \{1, 2\}, (u_i), (S_i) \rangle$, where $S_i = \mathbb{R}_+$ and, for any $(s_1, s_2) \in S_1 \times S_2$,

$$\begin{aligned} u_1(s_1, s_2) &= s_1(1 - (s_1 + s_2)), \\ u_2(s_1, s_2) &= s_2(1 - (s_1 + s_2)). \end{aligned}$$

- Here s_i is to be interpreted as quantity produced, and $1 - (s_1 + s_2)$ is the inverse demand function
- Taking the derivative gives the effect of a marginal increase in s_i on i 's payoff:

$$\frac{\partial u_i(s_i, s_j)}{\partial s_i} = 1 - s_j - 2s_i. \quad (1)$$

Example: Cournot model with linear demand (2)

- If (1) is positive (negative) under (s_i, s_j) , then marginally increasing (decreasing) s_i increases i 's payoff. If this holds continuously in the interval $[a, b]$ of i 's choices under s_j , then increasing s_i from a to b increases i 's payoff.
- By (1), $s_i = 1/2$ strictly dominates any $s_i > 1/2$, given that $s_j \geq 0$. Thus

$$S_i^1 = \left\{ s_i : 0 \leq s_i \leq \frac{1}{2} \right\}, i = 1, 2.$$

- By (1), $s_i = 1/2 - (1/2)^2$ strictly dominates any $s_i < 1/2 - (1/2)^2$, given that $0 \leq s_j \leq 1/2$. Thus

$$S_i^2 = \left\{ a_i : \frac{1}{2} - \left(\frac{1}{2} \right)^2 \leq a_i \leq \frac{1}{2} \right\}, i = 1, 2.$$

Example: Cournot model with linear demand (3)

- 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 $1/2 - (1/2)^2 \leq a_j \leq 1/2$. Thus

$$S_i^3 = \left\{ a_j : \frac{1}{2} - \left(\frac{1}{2}\right)^2 \leq a_j \leq \frac{1}{2} - \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 \right\}, i = 1, 2.$$

- Continuing this way for k (odd) steps, we get

$$S_i^k = \left\{ a_i : \begin{array}{c} \frac{1}{2} - \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 - \dots - \left(\frac{1}{2}\right)^{k-1} \\ \leq a_i \leq \\ \frac{1}{2} - \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 - \dots + \left(\frac{1}{2}\right)^k \end{array} \right\}.$$

Example: Cournot model with linear demand (4)

- 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.

Rationalizability

- Iterated strict dominance eliminates all the strategies that are dominated
- Perhaps we could be even more selective: eliminate all the strategies that are not *best responses* to a reasonable belief about the opponents strategy
- This leads to the concept of rationalizability
- But it turns out that this concept is (almost) equivalent to the concept of iterated strict dominance

- We say that a strategy of player i is rationalizable when it is a best response to a "reasonable" belief of i concerning the other players' actions
- By a "belief" of i , we mean a probability distribution over the other players' pure strategies:

$$\mu_i \in \Delta(S_{-i}),$$

where

$$S_{-i} = \times_{j \in \mathcal{I} \setminus i} S_j.$$

- There are two notions of rationalizability in the literature: either μ_i is a joint probability distribution allowing other players' actions to be correlated, or more restrictively, other player's actions are required to be independent
- Unlike MWG, we allow here correlation (this is sometimes called *correlated rationalizability*).

- What is a "reasonable" belief of i regarding other players' actions?
- Building on the notion of "common knowledge of rationality", beliefs should put positive weight only on those other players' strategies, which in turn can be rationalized
- Formally, this leads to a following definition (assume finite strategy sets).

Definition

The set of rationalizable strategies is the largest set $\times_{i \in \mathcal{I}} Z_i$, where $Z_i \subseteq S_i$, and each $s_i \in Z_i$ is a best-response to some belief $\mu_i \in \Delta(Z_{-i})$.

- To make the link to iterated strict dominance, define:

Definition

A strategy s_i is a **never-best response** if it is not a best response to any belief $\mu_i \in \Delta(S_{-i})$.

Theorem

A strategy s_i is a never-best response if and only if it is strictly dominated.

Proof of the theorem (1)

- It is clear that a strictly dominated strategy is never a best response. The challenge is to prove the converse, that a never-best response is strictly dominated.
- By contrapositive, we need to show that if strategy s_i is not strictly dominated, then it is a best-response given some belief $\mu_i \in \Delta(S_{-i})$.
- Let $\bar{u}_i(\sigma_i) := \{u_i(\sigma_i, s_{-i})\}_{s_{-i} \in S_{-i}}$ denote a vector, where each component is i 's payoff with mixed strategy σ_i , given a particular pure strategy profile s_{-i} for the other players.
- This vector contains i 's value for all possible combinations of pure strategies possible for the other players.

Proof of the theorem (2)

- Let N denote the number of elements of that vector so that $\bar{u}_i(\sigma_i) \in \mathbb{R}^N$. Given an arbitrary belief $\mu_i \in \Delta(S_{-i})$, we can then write the payoff for strategy σ_i as:

$$u(\sigma_i, \mu_i) := \mu_i \cdot \bar{u}_i(\sigma_i).$$

- Consider the set of such vectors over all σ_i :

$$U_i := \{\bar{u}_i(\sigma_i)\}_{\sigma_i \in \Sigma_i}.$$

- It is clear that U_i is a convex set.

Proof of the theorem (3)

- Assume that s_i is not strictly dominated. Let $U^+(s_i)$ denote the set of payoff vectors that strictly dominate $\bar{u}_i(s_i)$:

$$U^+(s_i) := \left\{ u \in \mathbb{R}^N : (u)_k > (\bar{u}_i(s_i))_k \text{ for all } k = 1, \dots, N \right\},$$

where $(\cdot)_k$ denotes the k^{th} component of a vector.

- $U^+(s_i)$ is a convex set, and since s_i is not strictly dominated, we have $U_i \cap U^+(s_i) = \emptyset$.
- By the separating hyperplane theorem, there exists some vector $\mu_i \in \mathbb{R}^N$, $\mu_i \neq 0$, such that

$$\mu_i \cdot (\bar{u}_i(\sigma_i) - \bar{u}_i(s_i)) \leq 0 \text{ for all } \sigma_i \in \Sigma_i \text{ and} \quad (2)$$

$$\mu_i \cdot (u - \bar{u}_i(s_i)) \geq 0 \text{ for all } u \in U^+(s_i). \quad (3)$$

Proof of the theorem (4)

- By (3), each component of μ_i must be positive.
- We can also normalize μ_i so that its components sum to one (without violating (2) or (3)), so that

$$\mu_i \in \Delta(S_{-i}).$$

Proof of the theorem (5)

- Equation (2) can now be written as

$$\mu_i \cdot \bar{u}_i(s_i) \geq \mu_i \cdot \bar{u}_i(\sigma_i) \text{ for all } \sigma_i \in \Sigma_i,$$

or

$$u(s_i, \mu_i) \geq u(\sigma_i, \mu_i) \text{ for all } \sigma_i \in \Sigma_i,$$

so that s_i is a best response to belief $\mu_i \in \Delta(S_{-i})$.

Rationalizable strategies

- Given this result, the process of iteratively deleting those strategies that are not best responses to any belief on the other players' remaining strategies is equivalent to the process of deleting strictly dominated strategies. Therefore, we have the following result:

Theorem

The set of pure strategies that survive the elimination of strictly dominated strategies is the same as the set of rationalizable strategies.

Remarks:

- Note: our definition of "never-best response" considers arbitrary belief $\mu_i \in \Delta(S_{-i})$ that allows i to believe that other players' actions are correlated
- If correlation not allowed, then the equivalence between "never-best response" and "strictly dominated" breaks down with more than two players: there are strategies that are never best responses to independent randomizations of the other players, yet they are not strictly dominated
- Hence, the alternative notion of "rationalizability" (that rules out correlation) is somewhat stronger than iterated strict dominance
- But this difference is not relevant in two-player games (because correlation between other players strategies is not relevant)

Discussion

- Rationalizability is the ultimate implication of common knowledge of rationality in games
- But it makes generally weak predictions. In many interesting games it does not imply anything. For example, consider the "Battle of sexes" game:

	<i>Ballet</i>	<i>Football</i>
<i>Ballet</i>	2, 1	0, 0
<i>Football</i>	0, 0	1, 2

- Rationalizability allows all outcomes. For example, players could choose (F, B) : F is optimal to player 1 who expects 2 to play F , and B is optimal to player 2 who expects 1 to play B .
- A way forward: require expectations to be mutually correct
 - Nash equilibrium

Nash equilibrium

- Rationalizability requires that each player's strategy is a best response to a reasonable conjecture on other player's play
 - *optimality principle*
- *Nash equilibrium* is a more stringent condition on strategic behavior.
- It requires that players play a best response against a *correct* belief of each other's play.
 - *consistency principle*

Nash equilibrium

Definition

A pure strategy profile $s = (s_1, \dots, s_I)$ constitutes a Nash equilibrium if for every $i \in \mathcal{I}$,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

for all $s'_i \in S_i$.

- Equivalently, s constitutes a Nash equilibrium if s_i is a best response against s_{-i} for all i .

Existence?

- Consider the game of Matching Pennies (or think about familiar Rock-Paper-Scissors game):

		2	
		<i>H</i>	<i>T</i>
1	<i>H</i>	1, -1	-1, 1
	<i>T</i>	-1, 1	1, -1

- Clearly, whenever player i chooses best response to j , j wants to change. There is no rest point for the best-response dynamics.
- Hence, there is no pure strategy Nash equilibrium

Nash equilibrium in mixed strategies

- This non-existence problem is avoided if we allow mixed strategies

Definition

A mixed strategy profile $\sigma \in \Sigma$ constitutes a Nash equilibrium if for every $i \in \mathcal{I}$,

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i})$$

for all $s_i \in S_i$.

- It is easy to check that playing H and T with prob. $1/2$ constitutes a Nash equilibrium of the matching pennies game
- Whenever σ is a Nash equilibrium, each player i is indifferent between all s_i for which $\sigma(s_i) > 0$. This is the key to solving for mixed strategy equilibrium.

Discussion

- A Nash equilibrium strategy is a best response, so it is rationalizable
- Hence, if there is a unique rationalizable strategy profile, then this profile must be a Nash equilibrium
- Obviously also: dominance solvability implies Nash
- An attractive feature of Nash equilibrium is that if players agree on playing Nash, then no player has an incentive to deviate from agreement
- Hence, Nash equilibrium can be seen as a potential outcome of preplay communication
- Keep in mind that Nash equilibrium can be seriously inefficient (Prisoner's dilemma)

Interpretations of mixed strategy equilibrium

- Do people really "randomize" their actions? Or should we interpret the mixed strategy equilibrium in some other way?
- There are various interpretations:
 1. Mixed strategies as objects of choice
 - This is the straightforward interpretation: people just randomize
 2. Mixed strategy equilibrium as a steady state
 - Players interact in an environment where similar situation repeats itself, without any strategic link between plays
 - Players know the frequencies with which actions were taken in the past

Interpretations of mixed strategy equilibrium

3. Mixed strategies as pure strategies in a perturbed game

- Players' preferences are subject to small perturbations
- Exact preferences are private information
- Mixed strategy equilibrium as the limit of pure strategy equilibrium of the perturbed game as perturbation vanishes
- This is the purification argument by Harsanyi (1973)

4. Mixed strategies as beliefs

- Think of σ as a belief system such that σ_i is the common belief of all the players of i 's actions
- Each action in the support of σ_i is best-response to σ_{-i}
- Each player chooses just one action
- Equilibrium is a steady state in the players' beliefs, not in their actions

Existence

- There are many existence results that guarantee existence of Nash equilibrium under varying conditions
- The best known applies to finite games, and was proved by Nash (1950):

Theorem

Finite games, i.e., games with I players and finite strategy sets S_i , $i = 1, \dots, I$, have a mixed strategy Nash equilibrium.

- The proof relies on the upper-hemicontinuity of the players' best-response correspondences, and the utilization of Kakutani's fixed point theorem
- See MWG Appendix of Ch. 8, or Mailath section 4.1 for proof (and mathematical appendix of MWG for upper-hemicontinuity)

Existence

- In many applications, it is more natural to model strategy space as a continuum
- Think about, e.g., Cournot oligopoly
- There is then no general existence theorem (it is easy to construct games without Nash equilibria)
- The simplest existence theorem assumes quasi-concave utilities:

Theorem

Assume that S_i are nonempty compact convex subsets of an Euclidean space, $u_i : S \rightarrow \mathbb{R}$ is continuous for all i and quasiconcave in s_i for all i . Then the game has a Nash equilibrium in pure strategies.

- Again, see MWG Appendix of Ch. 8 for the proof.
- In fact, Nash's theorem (previous theorem) is a special case of this
- Many other theorems apply to various situations where continuity and/or quasiconcavity fail
- There are situations though where existence fails

Multiplicity

- A more serious concern for game theory is the multiplicity of equilibria. Consider Battle of sexes

	<i>Ballet</i>	<i>Football</i>
<i>Ballet</i>	2, 1	0, 0
<i>Football</i>	0, 0	1, 2

or, Hawk-Dove

	<i>Dove</i>	<i>Hawk</i>
<i>Dove</i>	3, 3	1, 4
<i>Hawk</i>	4, 1	0, 0

- Both of these games have two pure strategy equilibria and one mixed strategy equilibrium (can you see this?)

Multiplicity

- Sometimes we can rule out an equilibrium using a **refinement** of the Nash Equilibrium concept
- For example, consider the following game:

	L	R
T	1, 1	0, 0
B	0, 0	0, 0

- Is BR a natural outcome?

Trembling hand perfect equilibrium

Definition

An equilibrium σ of a finite normal form game is a normal form *trembling hand perfect* equilibrium if there exists a sequence $\{\sigma^k\}_{k=1}^{\infty}$ of completely mixed strategy profiles such that

- $\sigma^k \rightarrow \sigma$
 - σ_i is a best reply to σ_{-i}^k for all k
- This is an example of an equilibrium refinement that captures Robustness to small mistakes
 - Rules out playing weakly dominated strategies
 - Always exists in a finite game
 - Rules out BR in the previous example

Multiplicity

- Sometimes we can use other criteria to *select* amongst many equilibria
- For example, equilibria may be pareto ranked
- Consider stag-hunt:

	A	B
A	2, 2	0, 1
B	1, 0	1, 1

- It can be argued that preplay communication helps to settle on pareto dominant equilibrium (A, A)

Multiplicity

- But even this might not be obvious. Consider:

	<i>A</i>	<i>B</i>
<i>A</i>	9, 9	0, 8
<i>B</i>	8, 0	7, 7

- Now playing *A* seems a bit shaky.. (what if the other player still chooses *B*?)
- Moreover, with preplay communication, players have an incentive to convince the other player that *A* will be played, even if they plan to play *B*. Is preplay communication credible?
- We conclude that there is no generally applicable answer for selecting among multiple equilibria

Correlated equilibrium

- Consider once again Battle of Sexes example
- There is a unique symmetric equilibrium in mixed strategies: each player takes her favourite action with a certain probability (compute this)
- But suppose that the players have a public randomization device (a coin for example). Let both players take the following strategy: go to ballet if heads, and to football if tails.
- Exercise: Show that this is an "equilibrium" and gives a better payoff to both players than the symmetric mixed strategy equilibrium.

- A generalization of this idea is called correlated equilibrium (see Osborne-Rubinstein Ch. 3.3, Fudenberg-Tirole Ch. 2.2, or Myerson Ch. 6 for more details)
- Correlated equilibrium may be interpreted as a solution concept that implicitly accounts for communication

Zero-sum games

- Let us end with a few words about a special class of games: zero-sum games
- A two-player game is a zero sum game if $u_1(s) = -u_2(s)$ for all $s \in S$.
- This of course implies that $u_1(\sigma) = -u_2(\sigma)$ for all σ .
- Matching pennies is a zero-sum game
- Zero-sum games are the most "competitive" games: maximizing ones payoff is equivalent to minimizing "opponent"'s payoff. There is no room for cooperation (should tennis players cooperate in Wimbledon final?)

- What is the largest payoff that player 1 can guarantee herself?
This is obtained by choosing

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2).$$

- Similarly for player 2:

$$\max_{\sigma_2 \in \Sigma_2} \min_{\sigma_1 \in \Sigma_1} u_2(\sigma_1, \sigma_2).$$

- But, because $u_1 = -u_2$, this is equivalent to

$$\min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2).$$

- Famous minmax theorem by von Neumann (1928) states that

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2).$$

- This maxmin value must give the payoff of a player in any Nash equilibrium (can you see why?)
- See e.g. Myerson Ch. 3.8 for more details.